arXiv:1704.02479v1 [stat.ME] 8 Apr 2017

Informed Bayesian T-Tests

Quentin F. Gronau*
Alexander Ly
and
Eric-Jan Wagenmakers

Department of Psychology, Psychological Methods,
University of Amsterdam, The Netherlands

*Correspondence concerning this article should be addressed to: Quentin F. Gronau, University of Amsterdam, Nieuwe Achtergracht 129 B, 1018 WT Amsterdam, The Netherlands. E-mail may be sent to Quentin.F.Gronau@gmail.com. This research was supported by a Netherlands Organisation for Scientific Research (NWO) grant to QFG (406.16.528) and a European Research Council (ERC) grant to EJW (283876). The authors would like to express their gratitude to Suzanne Oosterwijk for taking part in the prior elicitation effort. The authors would furthermore like to thank Michèle Nuijten for her help with obtaining the frequency of tests reported in Nuijten et al. (2016). R code can be found on the Open Science Framework: `https://osf.io/37vch/`.

Abstract

Across the empirical sciences, few statistical procedures rival the popularity of the frequentist t-test. In contrast, the Bayesian versions of the t-test have languished in obscurity. In recent years, however, the theoretical and practical advantages of the Bayesian t-test have become increasingly apparent. Developed by Harold Jeffreys, the default Bayesian t-test assigns a zero-centered Cauchy prior distribution to the effect size parameter under the alternative hypothesis. This specification is relatively restrictive, and in scientific applications the researcher may possess expert knowledge that allows the formulation of an informed prior distribution that is centered away from zero. In order to facilitate a more informative Bayesian analysis of the comparison between two means we extend Harold Jeffreys's t-tests. This extension allows researchers to assign a shifted and scaled t prior distribution to the effect size parameter. We present expressions for the marginal posterior distribution of the effect size parameter and for the Bayes factor test; three real-data examples highlight the practical value of the approach.

Keywords: informed hypothesis test, posterior distribution, Bayes factor

1 INTRODUCTION

The t-test is designed to assess whether two means differ from one another. The question is fundamental, and consequently the t-test has grown to be an inferential workhorse of the empirical sciences. The popularity of the t-test is underscored by considering the p-values published in eight major psychology journals from 1985 until 2013 (Nuijten et al., 2016); out of a total of 258,105 p-values, 26% tested the significance of a t statistic. For comparison, 4% of those p-values tested an r statistic, 4% a z statistic, 9% a χ^2 statistic, and 57% an F statistic. Similarly, Wetzels et al. (2011) found 855 t-tests reported in 252 psychology articles, for an average of about 3.4 t-tests per article.

The popularity of the t-test has always concerned its classical or frequentist version. This article concerns the Bayesian version originally developed by Jeffreys (1948). Jeffreys's Bayesian t-test quantifies the predictive adequacy of two competing hypotheses: the null hypothesis $\mathcal{H}_0$ which states that the effect size δ is equal to zero (i.e., $\mathcal{H}_0 : \delta = 0$) versus the alternative hypothesis $\mathcal{H}_1$ which assigns δ a zero-centered Cauchy prior distribution with scale parameter r (i.e., $\mathcal{H}_1 : \delta \sim \text{Cauchy}(0, r)$). Jeffreys's default choice was to set $r = 1$ (i.e., the standard Cauchy distribution).

Jeffreys's Bayesian t-tests come with the well-known advantages of Bayesian inference: the ability to assess the relative plausibility of different values for δ by means of a posterior distribution; the ability to quantify evidence, both in favor of the null hypothesis and in favor of the alternative hypothesis; the ability to monitor the evidential flow as more observations become available; the ability to stop data collection whenever the evidence is compelling (Berger and Berry, 1988; Rouder, 2014; Savage, 1961); and the ability to provide composite estimates and predictions that take into account model uncertainty (Haldane, 1932; Hoeting et al., 1999). The fact that Jeffreys's t-tests have remained relatively unpop-

ular among practitioners can be explained by a combination of factors: Jeffreys's writing is difficult to parse, Bayesian inference is generally excluded from the standard statistical curriculum offered in the empirical sciences, and the Bayesian t-tests were not available in mainstream statistical software packages. Recently, Rouder et al. (2009) have brought the Bayesian t-test into the limelight, with further impetus from the increased need for replication research to quantify evidence in favor of the null hypothesis (Marsman et al., 2017), and the fact that the test was subsequently implemented in user-friendly software packages (i.e., JASP Team, 2017; Morey and Rouder, 2015).

Jeffreys's Bayesian t-tests were meant to be "objective" in the sense of being applicable across a large range of different fields and phenomena. In Jeffreys's vision, possibly inspired by his experience as a physicist, the null hypothesis $\mathcal{H}_0$ represents an invariance or general law (i.e., $\delta = 0$), whereas the alternative hypothesis $\mathcal{H}_1$ relaxes the restriction imposed by $\mathcal{H}_0$ and reflects the intuition that the effect is likely to be small. Moreover, by centering the Cauchy prior on zero, Jeffreys avoided the complexity and context-dependency of having to specify another parameter (i.e., the most likely value of δ under $\mathcal{H}_1$). Nevertheless, a more informed specification of the prior distribution on δ, when possible, may increase the informativeness of the inference.

Here we present an extension of Jeffreys's default t-test that allows researchers to incorporate expert knowledge into the prior specification of the effect size parameter δ. Specifically, we consider two families of prior distributions for δ: the family of shifted and scaled t distributions (which includes Jeffreys's Cauchy prior as a special case) and the family of shifted and scaled normal distributions. For both families we derive the marginal posterior distribution of δ and the Bayes factor. For the normal family the solutions are completely analytic; for the t family the solutions contain a one-dimensional integral that can easily be evaluated numerically.

2 JEFFREYS'S DEFAULT BAYES FACTOR

Jeffreys developed Bayesian hypothesis tests for scenarios where the null hypothesis $\mathcal{H}_0$ represents a general law and the competing alternative hypothesis $\mathcal{H}_1$ relaxes the restriction imposed by $\mathcal{H}_0$. For instance, in the one-sample t-test setting, $\mathcal{H}_0$ states that the population mean is equal to zero whereas the alternative hypothesis allows for a non-zero population mean.

To quantify the evidence that the data provide for or against a general law, Jeffreys (1961) developed a formal system of statistical inference with the following key equation (Wrinch and Jeffreys, 1921, p. 387):

$$\underbrace{\frac{p(\mathcal{H}_1 \mid \text{data})}{p(\mathcal{H}_0 \mid \text{data})}}_{\text{Posterior odds}} = \underbrace{\frac{p(\mathcal{H}_1)}{p(\mathcal{H}_0)}}_{\text{Prior odds}} \times \underbrace{\frac{p(\text{data} \mid \mathcal{H}_1)}{p(\text{data} \mid \mathcal{H}_0)}}_{\text{Bayes factor BF}_{10}}. \tag{1}$$

In his work, Jeffreys focused on the *Bayes factor* (Etz and Wagenmakers, 2017; Kass and Raftery, 1995; Ly et al., 2016b; Robert et al., 2009). The Bayes factor has an intuitive interpretation: when $\text{BF}_{10} = 10$ this indicates that the data are 10 times more likely under $\mathcal{H}_1$ than under $\mathcal{H}_0$; when $\text{BF}_{10} = .2$ this indicates that the data are 5 times more likely under $\mathcal{H}_0$ than under $\mathcal{H}_1$.

Let $\mathcal{H}_0$ be specified by a series of nuisance parameters ζ and, importantly, a parameter of interest that is fixed at a single value of particular interest, $\theta = \theta_0$. The alternative hypothesis $\mathcal{H}_1$ is specified using similar nuisance parameters ζ, but in addition $\mathcal{H}_1$ releases the restriction on θ. In order to obtain the Bayes factor, the model parameters are integrated out

$$\text{BF}_{10} = \frac{\int_\Theta \int_Z p(\text{data} \mid \theta, \zeta, \mathcal{H}_1)\, p(\theta, \zeta \mid \mathcal{H}_1)\, \mathrm{d}\zeta\, \mathrm{d}\theta}{\int_Z p(\text{data} \mid \theta = \theta_0, \zeta, \mathcal{H}_0)\, p(\zeta \mid \mathcal{H}_0)\, \mathrm{d}\zeta}, \tag{2}$$

showing that the Bayes factor can be regarded as the ratio of two weighted averages where the weights correspond to the prior distribution for the parameters. Consequently, the

choice of the prior distribution is crucial for the development of a Bayes factor hypothesis test.

Jeffreys's procedure for assigning prior distributions to the parameters exploits the fact that $\mathcal{H}_0$ is nested within $\mathcal{H}_1$ (obtained by setting $\theta = \theta_0$). For the nuisance parameters, Jeffreys used his famous translation invariant "Jeffreys's" prior given by $p(\zeta) \propto \sqrt{\det \mathcal{I}(\zeta)}$ where $\mathcal{I}(\zeta)$ denotes the Fisher information matrix. For the test-relevant parameter θ, Jeffreys based his prior choice on two desiderata. The first desideratum, *predictive matching*, states that the Bayes factor should be perfectly indifferent (i.e., $\text{BF}_{10} = 1$) in case the data are completely uninformative. The second desideratum, *information consistency*, states that the Bayes factor should provide infinite support for one of the two hypotheses in case the data are overwhelmingly informative (Bayarri et al., 2012; Jeffreys, 1942). In case of the t-test, these two desiderata led Jeffreys to a Cauchy prior distribution for the test-relevant effect size parameter (Ly et al., 2016b).

Jeffreys's development of the Bayes factor test was mainly concerned with situations where there is little prior knowledge available. However, it seems that he was not completely opposed to including existing expert prior knowledge into the specification of the test-relevant prior distribution in case it is available (Jeffreys, 1961, p. 252):

> "In any of these cases it would be perfectly possible to give a form of $f(a)$ [i.e., prior distribution for the test-relevant parameter, in our notation $p(\theta)$] that would express the previous information satisfactorily, and consideration of the general argument of 5.0 will show that it would lead to common-sense results, but they would differ in scale. As we are aiming chiefly at a theory that can be used in the early stages of a subject, we shall not at present consider the last type of case"

As Jeffreys alludes to, there may be situations where researchers have strong prior knowl-

edge about the subject area. In this case, a procedure that admits a prior distribution for the test-relevant parameter which better reflects the prior state of knowledge might yield more informative tests of the researchers' hypotheses than tests based on the default prior.

Note that Jeffreys's default t-test assigns the crucial effect size parameter δ a zero-centered Cauchy distribution under $\mathcal{H}_1$; hence, under $\mathcal{H}_1$ the most likely effect size value is zero, a proposition that many researchers find problematic. Instead, these researchers may wish to center their prior distribution on the non-zero value for δ that reflects their knowledge-based expectation about the effects at hand. Below we derive the posterior distributions and Bayes factors for two families of prior distributions on δ that allow changes in shift and scale. We first consider the one-sample and paired samples t-test and then consider the independent samples t-test.

3 ONE-SAMPLE AND PAIRED SAMPLES T-TEST

3.1 Model

In a one-sample t-test, n observations are assumed to be drawn from a normal distribution, that is, $y_i \sim \mathcal{N}(\mu, \sigma^2)\, \forall i \in \{1, 2, \ldots, n\}$. Note that the paired samples t-test is a special case of this formalization, since a paired samples t-test is equivalent to a one-sample t-test on the difference scores. We follow Jeffreys and reparameterize the model in terms of the dimensionless effect size $\delta = \mu/\sigma$ (Jeffreys, 1961; Ly et al., 2016b; Rouder et al., 2009; Wetzels et al., 2009). The null hypothesis states that the effect size is zero (i.e., $\mathcal{H}_0 : \delta = 0$) whereas the alternative hypothesis states that the effect size is non-zero (i.e., $\mathcal{H}_1 : \delta \neq 0$).

3.2 Prior Distributions

Next we need to specify prior distributions for the parameters. For the parameter common to both models (i.e., σ^2), we use $p(\sigma^2) \propto \sigma^{-2}$, in line with Jeffreys's approach. For the test-relevant effect size δ, we depart from Jeffreys's default choice and consider an informed shifted and scaled t prior distribution, that is, $\delta \sim t(\mu_\delta, r, \kappa)$, where μ_δ corresponds to the location parameter, r to the scale parameter, and κ to the degrees of freedom. Note that this specification includes Jeffreys's standard Cauchy prior distribution as a special case (i.e., when $\kappa = 1$). Additionally, we present the results for an informed normal prior of the form $\delta \sim \mathcal{N}(\mu_\delta, g)$ where μ_δ corresponds to the mean and g to the variance of the normal distribution.

3.3 Posterior Distribution for Effect Size δ

After conducting a Bayesian t-test, the two quantities of primary interest are the posterior distribution for effect size δ under $\mathcal{H}_1$ and the Bayes factor (presented in the next section). For enhanced readability we suppress the conditioning on $\mathcal{H}_1$ and write the marginal posterior distribution for δ under $\mathcal{H}_1$ as follows:

$$p(\delta \mid y) = \frac{\int_0^\infty p(y \mid \delta, \sigma^2)\, p(\sigma^2) \mathrm{d}\sigma^2\, p(\delta)}{p(y)}. \tag{3}$$

When we use a t prior of the form $t(\mu_\delta, r, \kappa)$, we obtain

$$p(\delta \mid y) = \frac{\exp\left\{-\frac{n}{2}\delta^2\right\} \left(1 + \frac{t^2}{\nu}\right)^{-\frac{\nu+1}{2}} [C + D]\, \frac{1}{r} \frac{\Gamma\left(\frac{\kappa+1}{2}\right)}{\sqrt{\pi\kappa}\,\Gamma\left(\frac{\kappa}{2}\right)} \left[1 + \frac{\left(\frac{\delta - \mu_\delta}{r}\right)^2}{\kappa}\right]^{-\frac{\kappa+1}{2}}}{\int_0^\infty (1 + ng)^{-\frac{1}{2}} \exp\left\{-\frac{\mu_\delta^2}{2\left(\frac{1}{n} + g\right)}\right\} \left(1 + \frac{t^2}{\nu(1+ng)}\right)^{-\frac{\nu+1}{2}} [A + B]\, p(g) \mathrm{d}g}, \tag{4}$$

where $p(g)$ denotes an inverse-gamma distribution of the form $IG\left(\frac{\kappa}{2}, r^2\frac{\kappa}{2}\right)$ given by

$$p(g) = \frac{\left(r^2\frac{\kappa}{2}\right)^{\frac{\kappa}{2}}}{\Gamma(\frac{\kappa}{2})} g^{-\frac{\kappa}{2}-1} \exp\left(-\frac{r^2\kappa}{2g}\right), \tag{5}$$

and

$$A = \Gamma\left(\frac{\nu+1}{2}\right) {}_1F_1\left(\frac{\nu+1}{2}; \frac{1}{2}; \frac{\mu_\delta^2 t^2}{2\left(\frac{1}{n}+g\right)\left[(1+ng)\,\nu + t^2\right]}\right), \tag{6}$$

$$B = \frac{\mu_\delta t}{\sqrt{\frac{1}{2}\left(\frac{1}{n}+g\right)\left[(1+ng)\,\nu + t^2\right]}} \Gamma\left(\frac{\nu+2}{2}\right) \tag{7}$$

$$\times {}_1F_1\left(\frac{\nu+2}{2}; \frac{3}{2}; \frac{\mu_\delta^2 t^2}{2\left(\frac{1}{n}+g\right)\left[(1+ng)\,\nu + t^2\right]}\right), \tag{8}$$

$$C = \Gamma\left(\frac{\nu+1}{2}\right) {}_1F_1\left(\frac{\nu+1}{2}; \frac{1}{2}; \frac{nt^2\delta^2}{2(\nu+t^2)}\right), \tag{9}$$

$$D = t\delta\sqrt{\frac{2n}{\nu+t^2}} \Gamma\left(\frac{\nu+2}{2}\right) {}_1F_1\left(\frac{\nu+2}{2}; \frac{3}{2}; \frac{nt^2\delta^2}{2(\nu+t^2)}\right), \tag{10}$$

where ${}_1F_1$ denotes the confluent hypergeometric function. Furthermore, t corresponds to the one-sample t-value defined as $t = \sqrt{n}\bar{y}/s$ where s corresponds to the unbiased sample standard deviation and $\nu = n - 1$ corresponds to the degrees of freedom of the t-test. The integral with respect to g in the denominator is due to the t prior representation as a scale-mixture of normal distributions as, for instance, applied in the well-known mixture of g priors (e.g., Liang et al., 2008). This integral is independent of δ and thus only needs to be evaluated once, something which can easily be accomplished via numerical integration.

In order to derive this expression for the posterior distribution we made use of a lemma distilled from the Bateman project (Bateman et al., 1954; Ly et al., 2016a). Note that in order to calculate the posterior distribution, we only require the t-value and the sample size; this is convenient because it allows a computation of the posterior distribution of δ for

a t-test reported in an empirical article which usually presents the t-value and the sample size, but not the raw data.

The posterior distribution for a normal prior of the form $\delta \sim \mathcal{N}(\mu_\delta, g)$ is obtained by removing the integral over g with respect to $p(g)$ in the denominator and replacing the shifted and scaled t prior distribution in the numerator by the normal prior distribution which yields:

$$p(\delta \mid y) = \frac{\exp\left\{-\frac{n}{2}\delta^2\right\}\left(1+\frac{t^2}{\nu}\right)^{-\frac{\nu+1}{2}}[C+D]\,(2\pi g)^{-\frac{1}{2}}\exp\left\{-\frac{1}{2g}(\delta-\mu_\delta)^2\right\}}{(1+ng)^{-\frac{1}{2}}\exp\left\{-\frac{\mu_\delta^2}{2\left(\frac{1}{n}+g\right)}\right\}\left(1+\frac{t^2}{\nu(1+ng)}\right)^{-\frac{\nu+1}{2}}[A+B]}, \tag{11}$$

where A, B, C, and D are defined as before. Detailed derivations can be found in the appendix.

3.4 Bayes Factor

In this section we present the Bayes factor for the informed generalization of Jeffreys's one-sample t-test. The Bayes factor BF_{10} is given by the ratio of the *marginal likelihoods* of $\mathcal{H}_1$ and $\mathcal{H}_0$. The marginal likelihoods are obtained by integrating out the model parameters with respect to their prior distributions:

$$\text{BF}_{10} = \frac{\int_0^\infty \int_{-\infty}^\infty p(y \mid, \delta, \sigma^2)\, p(\delta)\, p(\sigma^2) \mathrm{d}\delta \mathrm{d}\sigma^2}{\int_0^\infty p(y \mid, \delta = 0, \sigma^2)\, p(\sigma^2) \mathrm{d}\sigma^2}. \tag{12}$$

The Bayes factor for a shifted and scaled t prior on δ is given by:

$$\text{BF}_{10} = \frac{\int_0^\infty (1+ng)^{-\frac{1}{2}}\exp\left\{-\frac{\mu_\delta^2}{2\left(\frac{1}{n}+g\right)}\right\}\left(1+\frac{t^2}{\nu(1+ng)}\right)^{-\frac{\nu+1}{2}}[A+B]\,p(g)\mathrm{d}g}{\Gamma\left(\frac{\nu+1}{2}\right)\left(1+\frac{t^2}{\nu}\right)^{-\frac{\nu+1}{2}}}, \tag{13}$$

the Bayes factor for a normal prior on δ is given by

$$\text{BF}_{10} = \frac{(1+ng)^{-\frac{1}{2}} \exp\left\{ -\frac{\mu_\delta^2}{2\left(\frac{1}{n}+g\right)} \right\} \left(1 + \frac{t^2}{\nu(1+ng)}\right)^{-\frac{\nu+1}{2}} [A+B]}{\Gamma\left(\frac{\nu+1}{2}\right)\left(1+\frac{t^2}{\nu}\right)^{-\frac{\nu+1}{2}}}, \tag{14}$$

where A, B, ν, t, and $p(g)$ are defined as before. Note that, as for the posterior distribution of δ, the only information needed to compute this Bayes factor is the t-value and the sample size n. A detailed derivation is presented in the appendix.

3.5 Example I: The Crowd Within Effect

As an example application, consider the so-called *crowd within effect*. When people are asked to provide quantitative judgments, it has long been known that the average of the estimates across persons tends to be more accurate than the individual judgments; this is called the *wisdom of the crowd effect* (Galton, 1907; Surowiecki, 2004). Recently, Vul and Pashler (2008) showed that a similar effect may occur when a single individual is asked to provide a quantitative judgment twice: the average of two successive assessments was more accurate than the first or second assessment alone. Due to the similarity to the wisdom of the crowd effect, Vul and Pashler (2008) termed this phenomenon the *crowd within effect*.

This surprising effect was successfully replicated by Steegen et al. (2014). Here, we present an informed reanalysis of the replication study where we use our prior knowledge from the original study to specify the effect size prior distribution for the replication experiment, following Lindley's adage "today's posterior is tomorrow's prior" (Lindley, 1972).

The original experiment by Vul and Pashler featured an immediate and a delayed condition. In the immediate condition, participants provided judgments to eight questions such as "What percentage of the world's airports are in the United States?" and they were asked

to make a second guess for each of the questions immediately after they had completed the questionnaire. In the delayed condition, participants provided the second judgments three weeks after completing the questionnaire. For our reanalysis, we focus on the results for the delayed condition. Furthermore, we only consider the comparison of the averaged estimate to the first guess. In the original experiment, a classical paired t-test indicated that the error of the average was smaller than the error of the first guess: $t(172) = 6.22, p < .001$. Steegen et al. successfully replicated this finding: $t(139) = 4.02, p < .001$. For our reanalysis, we proceeded as follows (Verhagen and Wagenmakers, 2014):

1. We analyzed the original experiment using a zero-centered Cauchy prior for the effect size as proposed by Jeffreys (1948). However, instead of using the standard Cauchy distribution that was Jeffreys's default choice, we set the scale parameter of the Cauchy distribution to $1/\sqrt{2} \approx 0.707$, the present default choice in the field of psychology. This Cauchy prior is equivalent to a $t\left(0, 1/\sqrt{2}, 1\right)$ prior distribution for the effect size parameter δ.

2. Using the default Cauchy prior distribution, the paired t-test reported in the original article yields a Bayes factor of $\text{BF}_{10} = 2,483,125$ indicating overwhelming evidence for $\mathcal{H}_1 : \delta \sim \text{Cauchy}(0, 1/\sqrt{2})$ over $\mathcal{H}_0 : \delta = 0$. This result can be obtained using the equations just presented, but the analysis is also easily conducted using the `BayesFactor` R package (Morey and Rouder, 2015) which also produces samples from the posterior distribution for the effect size parameter δ. We fitted a t distribution to these posterior samples which yielded the following distribution: $t(0.465, 0.078, 41.478)$.

3. Next, the posterior distribution for δ served as the prior distribution for the analysis of the replication experiment. Specifically, we contrasted the skeptic's $\mathcal{H}_0 : \delta = 0$

to the proponent's $\mathcal{H}_F : \delta \sim t(0.465, 0.078, 41.478)$. We obtain a Bayes factor of $\text{BF}_{F0} = 901.5$ indicating that the data from Steegen et al. are about 900 times more likely under $\mathcal{H}_F$ than under $\mathcal{H}_0$. As a comparison, we also conducted the default analysis of the replication experiment (i.e., using again the zero-centered Cauchy prior with scale parameter $1/\sqrt{2}$ instead of the informed t prior distribution) which yielded a Bayes factor of $\text{BF}_{10} = 170.2$.

Figure 1 displays the informed prior distribution (dotted line), the posterior distribution after observing the replication experiment (solid line), a 95% posterior credible interval, the posterior median, and the Bayes factor. The posterior distribution is located near the informed prior distribution, albeit centered on a slightly smaller effect size value.

Figure 2 displays the results for the default Cauchy prior distribution with a scale parameter of $1/\sqrt{2}$ (dotted line). Compared to the posterior distribution based on the informed t prior, the posterior distribution based on the default Cauchy prior is wider and shifted towards smaller effect size values. This reflects the influence of the prior distribution which, in case of the informed t prior, "pulls" the posterior towards larger values.

Another observation is that the Bayes factor in favor of the informed alternative hypothesis is larger than the Bayes factor in favor of the default alternative hypothesis. This can be explained by interpreting the Bayes factor as a measure of the predictive success of two competing hypotheses. The informed alternative hypothesis makes riskier predictions than the default alternative hypothesis, but since the results are consistent with these predictions, the informed hypothesis is rewarded more. In fact, we can obtain the Bayes factor that compares the informed alternative hypothesis to the default alternative hypothesis by transitivity:

$$\text{BF}_{F1} = \underbrace{\frac{p(y \mid \mathcal{H}_F)}{p(y \mid \mathcal{H}_0)}}_{\text{BF}_{F0}} \times \underbrace{\frac{p(y \mid \mathcal{H}_0)}{p(y \mid \mathcal{H}_1)}}_{\text{BF}_{01}} = \frac{\text{BF}_{F0}}{\text{BF}_{10}}. \tag{15}$$

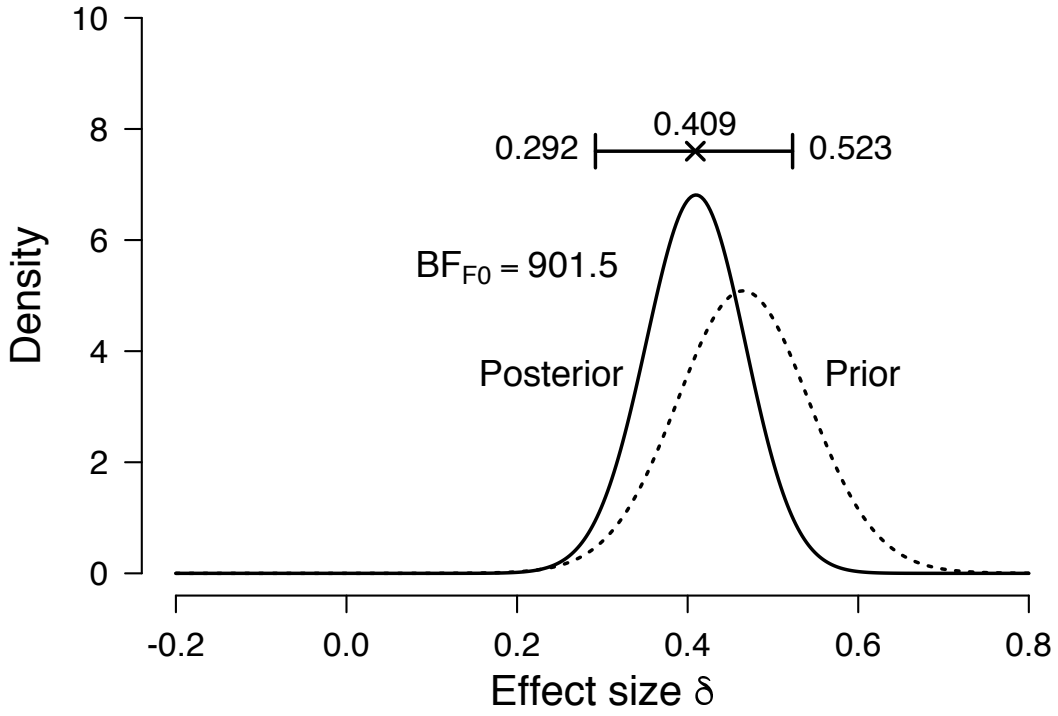

Figure 1: Results of the informed reanalysis of the crowd within effect replication study by Steegen et al. (2014). The dotted line corresponds to the informed $t(0.465, 0.078, 41.478)$ prior distribution based on our knowledge from the original study by Vul and Pashler (2008). The solid line corresponds to the posterior distribution, with a 95% credible interval and the posterior median displayed on top. The Bayes factor in favor of the informed alternative hypothesis over the null hypothesis equals $\text{BF}_{F0} = 901.5$. Figure available at `https://tinyurl.com/kzrhad6` under CC license `https://creativecommons.org/licenses/by/2.0/`.

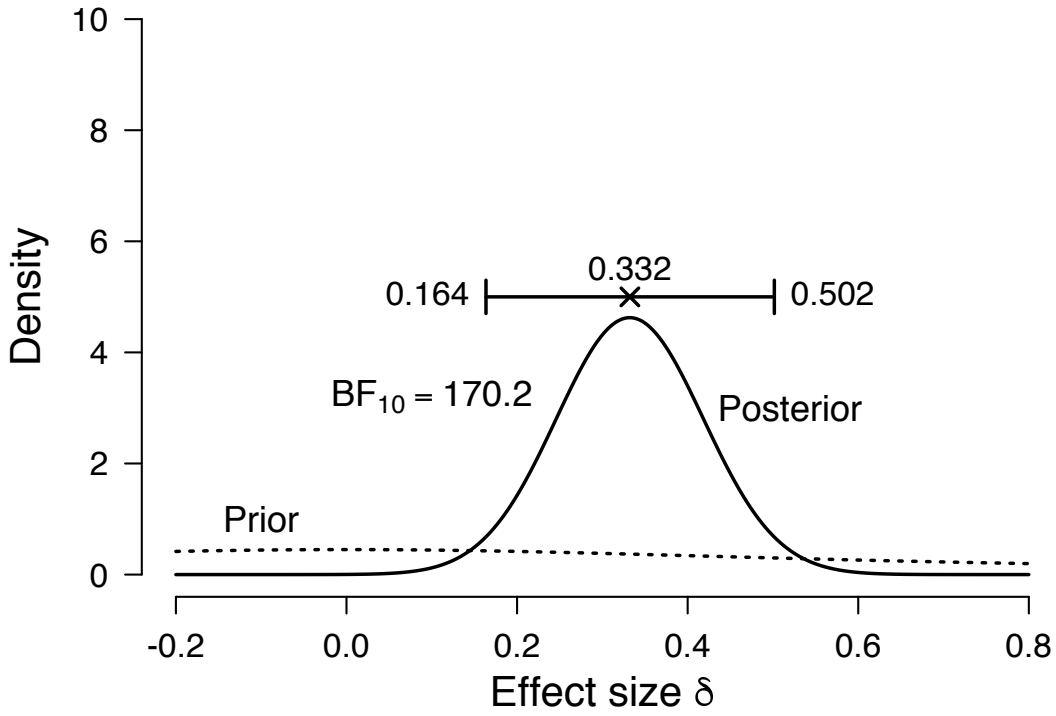

Figure 2: Results of the default analysis of the crowd within effect replication study by Steegen et al. (2014). The dotted line corresponds to the default Cauchy prior distribution with scale parameter $1/\sqrt{2}$. The solid line corresponds to the posterior distribution, with a 95% credible interval and the posterior median displayed on top. The Bayes factor in favor of the default alternative hypothesis equals $\text{BF}_{10} = 170.2$. Figure available at `https://tinyurl.com/n7tkm9y` under CC license `https://creativecommons.org/licenses/by/2.0/`.

Hence, the Bayes factor in favor of the informed alternative hypothesis over the default alternative hypothesis equals $\text{BF}_{F1} = 901.5/170.2 \approx 5.3$.

4 UNPAIRED SAMPLES T-TEST

4.1 Model

We extend our informed approach to the unpaired/independent samples version of Jeffreys's t-test. We assume that n_y observations are drawn from a normal distribution of the form $y_i \sim \mathcal{N}(\mu - \frac{\alpha}{2}, \sigma^2) \, \forall i \in \{1, 2, \ldots, n_y\}$ in the first group, and n_x observations are drawn from a normal distribution of the form $x_j \sim \mathcal{N}(\mu + \frac{\alpha}{2}, \sigma^2) \, \forall j \in \{1, 2, \ldots, n_x\}$ in the second group. Similar to the one-sample t-test, we reparameterize the model in terms of the effect size $\delta = \alpha/\sigma$. As before, the null hypothesis $\mathcal{H}_0 : \delta = 0$ is pitted against the alternative hypothesis $\mathcal{H}_1 : \delta \neq 0$.

4.2 Prior Distributions

The unpaired samples t-test features two nuisance parameters (i.e., the variance σ^2 and the grand mean μ). In line with Jeffreys's approach, we assign these parameters the translation invariant prior $p(\mu, \sigma^2) \propto \sigma^{-2}$. For effect size δ, we again consider an informed shifted and scaled t distribution, that is $\delta \sim t(\mu_\delta, r, \kappa)$, and an informed normal prior, that is $\delta \sim \mathcal{N}(\mu_\delta, g)$.

4.3 Posterior Distribution for Effect Size δ and Bayes Factor

Conveniently, for the independent samples t-test, the expressions for the posterior distribution and the Bayes factor are obtained by adjusting the respective expressions for the

one-sample/paired samples t-test (i.e., Eq. 4 & Eq. 11 and Eq. 13 & Eq. 14) in the following way:

1. Replace n by an "effective" sample size n_{eff} defined as $n_{\text{eff}} = \frac{n_y n_x}{n_y + n_x}$ (Rouder et al., 2009; Ly et al., 2016b).

2. Replace the one-sample/paired samples degrees of freedom $\nu = n - 1$ by the independent samples degrees of freedom $\nu = n_y + n_x - 2$.

3. Replace the one-sample/paired samples t-test t-value by the independent samples t-value defined as

$$t = \frac{\bar{x} - \bar{y}}{\sqrt{\frac{(n_x - 1)s_x^2 + (n_y - 1)s_y^2}{n_y + n_x - 2} \times \left(\frac{1}{n_x} + \frac{1}{n_y}\right)}}. \tag{16}$$

A detailed derivation is presented in the appendix.

4.4 Example II: The Facial Feedback Hypothesis

The *facial feedback hypothesis* states that affective responses can be influenced by one's facial expression even when that facial expression is not the result of an emotional experience. In a seminal study, Strack et al. (1988) found that participants who held a pen between their teeth (inducing a facial expression similar to a smile) rated cartoons as more funny than participants who held a pen with their lips (inducing a facial expression similar to a pout).

In a recently published Registered Replication Report (Wagenmakers et al., 2016), 17 labs worldwide attempted to replicate this finding using a preregistered and independently vetted protocol. A random-effects meta-analysis yielded an estimate of the mean difference between the "smile" and "pout" condition equal to 0.03 [95% CI: -0.11, 0.16]. Furthermore,

one-sided default Bayesian unpaired t-tests (using a Cauchy prior scale of $1/\sqrt{2}$) revealed that for all 17 studies, the Bayes factor indicated evidence in favor of the null hypothesis and for 13 out of the 17 studies, the Bayes factor in favor of the null was larger than 3. Overall, the authors concluded that "the results were inconsistent with the original result" (Wagenmakers et al., 2016, p. 924).

Here we present an informed reanalysis of the replication studies based on a prior elicitation effort with Dr. Suzanne Oosterwijk, a social psychologist at the University of Amsterdam with considerable expertise in this domain.

4.4.1 Prior Elicitation

Before commencing the elicitation process, we asked our expert to ignore the knowledge about the failed replication of Strack et al. (1988). Next, we stressed that the goal of the elicitation effort was to obtain an informed prior distribution for δ *under the alternative hypothesis*, that is, under the assumption that the effect is present. This was important in order to prevent unwittingly eliciting a prior that is a mixture between a point mass at zero and the distribution of interest. Then, we proceeded in steps of increasing sophistication. First, together with the expert we decided that the theory specified a direction, implying a one-sided hypothesis test. Next, we asked the expert to provide a value for the median of the effect size: this yielded a value of 0.35. Subsequently, we asked for values for the 33% and 66% percentile of the prior distribution for the effect size: this yielded values of 33%-tile = 0.25 and 66%-tile = 0.45. To finesse and validate the specified prior distribution we used the MATCH Uncertainty Elicitation Tool (`http://optics.eee.nottingham.ac.uk/match/uncertainty.php`), a web application that allows one to elicit probability distributions from experts (Morris et al., 2014). The complete elicitation effort took approximately one hour and resulted in a t distribution with location 0.350, scale 0.102, and 3 degrees of

freedom.

4.4.2 Reanalysis of the Oosterwijk Replication Study

Having elicited an informed prior distribution for δ under $\mathcal{H}_F$, we now turn to a detailed reanalysis of the facial feedback replication attempt from Dr. Oosterwijk's lab at the University of Amsterdam. Later, we summarize the results for all 17 replication attempts.

The alternative hypothesis is directional, that is, the teeth condition is predicted to result in relatively high funniness ratings, not relatively low funniness ratings. In order to respect the directional nature of the alternative hypothesis the two-sided informed t-test outlined above requires a correction. Specifically, the Bayes factor that compares an alternative hypothesis that only allows for positive effect size values to the null hypothesis can be computed via a simply identity that exploits the transitive nature of the Bayes factor (Morey and Wagenmakers, 2014):

$$\text{BF}_{+0} = \underbrace{\frac{p(y \mid \mathcal{H}_+)}{p(y \mid \mathcal{H}_1)}}_{\text{BF}_{+1}} \times \underbrace{\frac{p(y \mid \mathcal{H}_1)}{p(y \mid \mathcal{H}_0)}}_{\text{BF}_{10}} = \text{BF}_{+1} \times \text{BF}_{10}. \tag{17}$$

We already showed how to obtain BF_{10}, that is, the Bayes factor for the two-sided test of an informed alternative hypothesis; BF_{+1} can be obtained by simply dividing the posterior mass for δ larger than zero (obtained by numerically integrating Eq. 4) by the prior mass for δ larger than zero. The Bayes factor hypothesis test that we report will respect the directional nature of the facial feedback hypothesis and include the correction term from Eq. 17.

Figure 3 displays the results for the reanalysis of the data from the Oosterwijk lab. The prior and posterior distribution do not impose the directional constraint. The one-sided informed Bayes factor equals $\text{BF}_{0F_+} = 11.5$ indicating that the data are about eleven

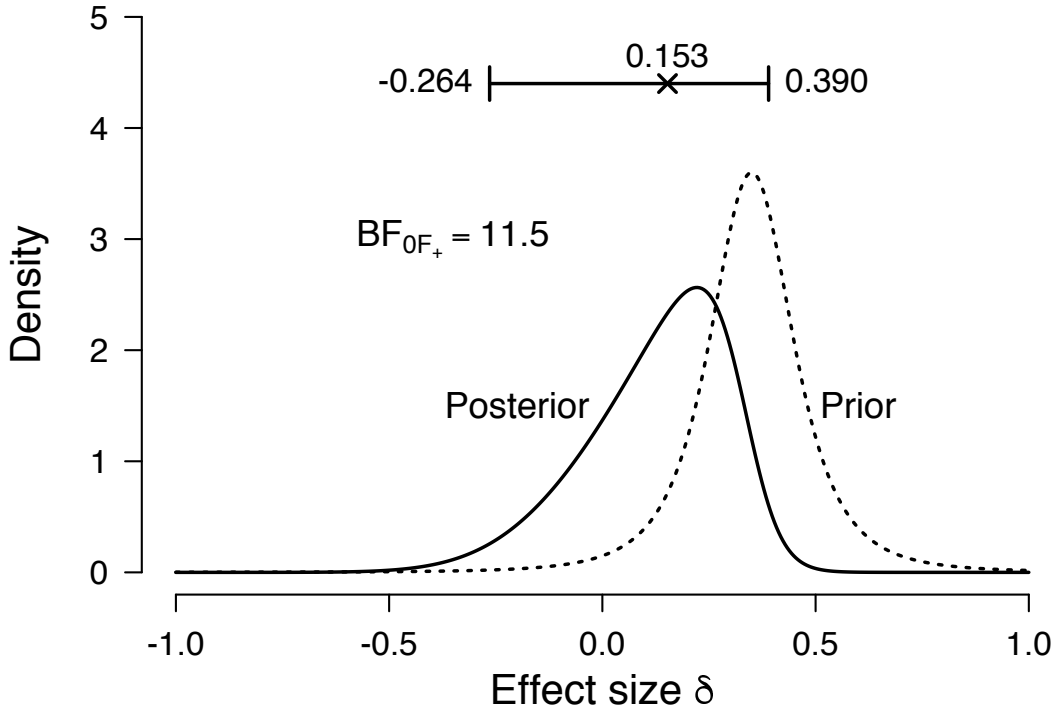

Figure 3: Results of an informed reanalysis of the facial feedback hypothesis replication data from the Oosterwijk lab. The dotted line corresponds to the elicited $t(0.350, 0.102, 3)$ prior distribution. The solid line corresponds to the associated posterior distribution, with a 95% credible interval and the posterior median displayed on top. The Bayes factor in favor of the null hypothesis over the one-sided informed alternative hypothesis equals $\text{BF}_{0F_+} = 11.5$. Figure available at `https://tinyurl.com/mk7uaxm` under CC license `https://creativecommons.org/licenses/by/2.0/`.

times more likely under the null hypothesis than under the one-sided informed alternative hypothesis.

For comparison, Figure 4 displays the results based on the default one-sided zero-centered Cauchy distribution with scale $1/\sqrt{2}$. The one-sided default Bayes factor equals $\text{BF}_{0+} = 8.7$, indicating that the data are about 9 times more likely under the null hypothesis than under the one-sided default alternative hypothesis. Hence, both the informed and the default Bayes factor yield the same qualitative conclusion, that is, considerable evidence for the null hypothesis. However, the unrestricted posterior distributions differ noticeably between the informed and the default analysis: the posterior median based on the informed prior specification is positive and equal to 0.153 (95% credible interval: $[-0.264, 0.390]$) whereas the posterior median based on the default prior distribution is equal to -0.152 (95% credible interval: $[-0.511, 0.200]$).

4.4.3 Reanalysis of all 17 Facial Feedback Replication Studies

Figure 5 displays the (nondirectional) posterior distributions under the informed $\mathcal{H}_F : \delta \sim t(0.350, 0.102, 3)$ and under the default $\mathcal{H}_1 : \delta \sim \text{Cauchy}(0, 1/\sqrt{2})$ for the reanalysis of all 17 replication studies. The posterior distributions for δ differ noticeably for the informed and the default prior specification. Under $\mathcal{H}_F$, the posterior distributions are shifted towards larger effect size values and the posteriors are more peaked than the ones corresponding to the default prior under $\mathcal{H}_1$.

Figure 6 displays the (directional) informed and default Bayes factors for the reanalysis of all 17 replication studies. Each dot corresponds to a study; the x-coordinate corresponds to the one-sided informed Bayes factor and the y-coordinate corresponds to the one-sided default Bayes factor as reported in Wagenmakers et al. (2016). The results are qualitatively similar. For Bayes factors smaller than about six, the one-sided default Bayes factor

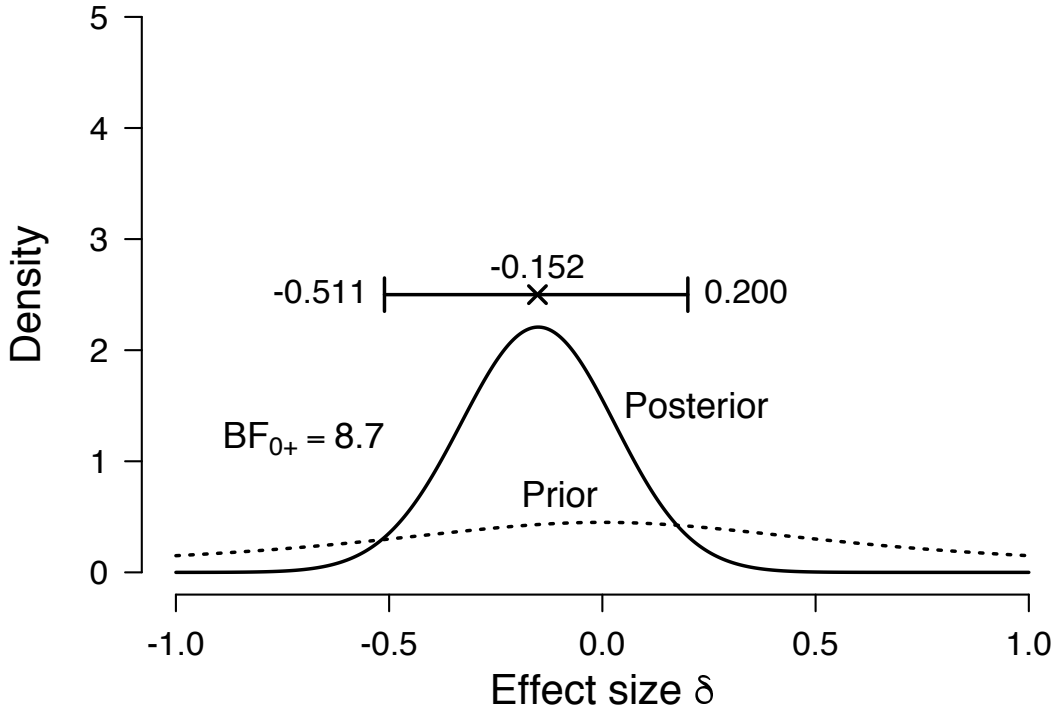

Figure 4: Results of the default analysis of the facial feedback hypothesis replication data from the Oosterwijk lab. The dotted line corresponds to the default Cauchy prior distribution with scale parameter $1/\sqrt{2}$. The solid line corresponds to the associated posterior distribution, with a 95% credible interval and the posterior median displayed on top. The Bayes factor in favor of the null hypothesis over the one-sided default alternative hypothesis equals $\text{BF}_{0+} = 8.7$. Figure available at `https://tinyurl.com/mgs28ob` under CC license `https://creativecommons.org/licenses/by/2.0/`.

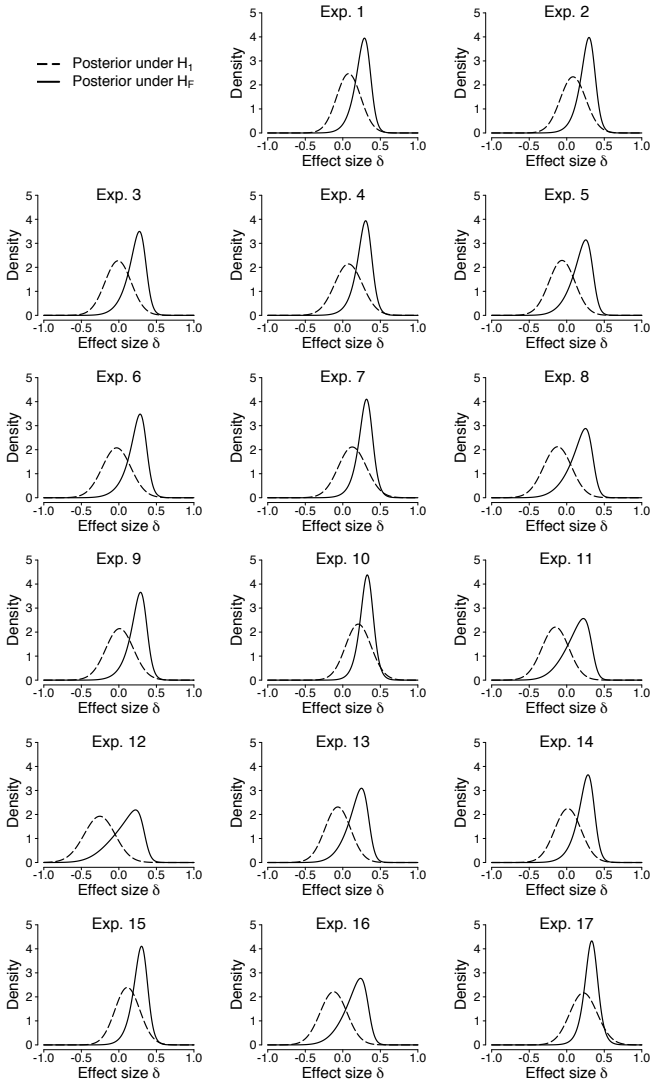

Figure 5: Comparison of the posterior distributions for δ under $\mathcal{H}_F: \ \delta \sim \ t(0.350, 0.102, 3)$ and $\mathcal{H}_1 : \delta \sim \ \text{Cauchy}(0, 1/\sqrt{2})$ for the facial feedback hypothesis replication data from the 17 labs. Figure available at `https://tinyurl.com/l7pxbno` under CC license `https://creativecommons.org/licenses/by/2.0/`.

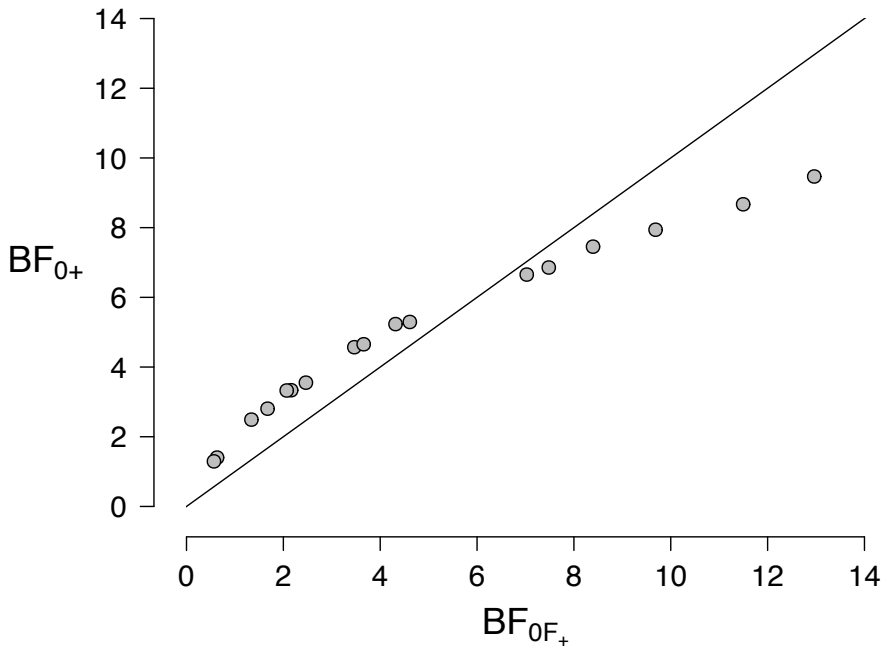

Figure 6: Comparison of the one-sided default and one-sided informed Bayes factor analysis of the facial feedback hypothesis replication data from the 17 labs. The x-coordinate corresponds to the one-sided informed Bayes factor and the y-coordinate corresponds to the one-sided default Bayes factor as reported in Wagenmakers et al. (2016). Figure available at `https://tinyurl.com/ke4489k` under CC license `https://creativecommons.org/licenses/by/2.0/`.

provides slightly more evidence for the null hypothesis than the one-sided informed Bayes factor. For Bayes factors larger than about six, this pattern is reversed.

In sum, for parameter estimation it matters whether the analysis is based on an informed prior distribution or a default prior distribution. In contrast, for hypothesis testing the results are relatively similar: both the informed Bayes factor and the default Bayes factor support the conclusion that the original study by Strack et al. (1988) could not be successfully replicated by the 17 labs involved in the replication attempt.

5 EXAMPLE III: REANALYSIS OF 593 T-TESTS

For our final example we reanalyze the significant subset of the 855 t-tests collected by Wetzels et al. (2011) from the 2007 issues of two popular psychology journals (i.e., *Psychonomic Bulletin & Review* and *Journal of Experimental Psychology: Learning, Memory, and Cognition*). The analysis of these 593 significant t-tests allows for an empirical investigation of the relation between the default Bayes factor and the informed Bayes factor.

For the informed Bayes factor, we adopt the expert prior from the previous example, that is, a $t(0.350, 0.102, 3)$ distribution for δ. This prior distribution was elicited within the context of the facial feedback hypothesis, but we believe it may serve as an informed prior for small-to-medium effect sizes across many areas of psychological research. For the default Bayes factor we again use the Cauchy$(0, 1/\sqrt{2})$ distribution for δ. Both Bayes factors are one-sided.

Figure 7 displays the one-sided informed Bayes factors (BF_{F_+0}) and the default one-sided Bayes factors (BF_{+0}) for the 593 significant t-tests collected by Wetzels et al. (2011). The two-sided p-values are indicated by different shapes and colors. The overall pattern indicates a positive correlation between the informed and default Bayes factor. However,

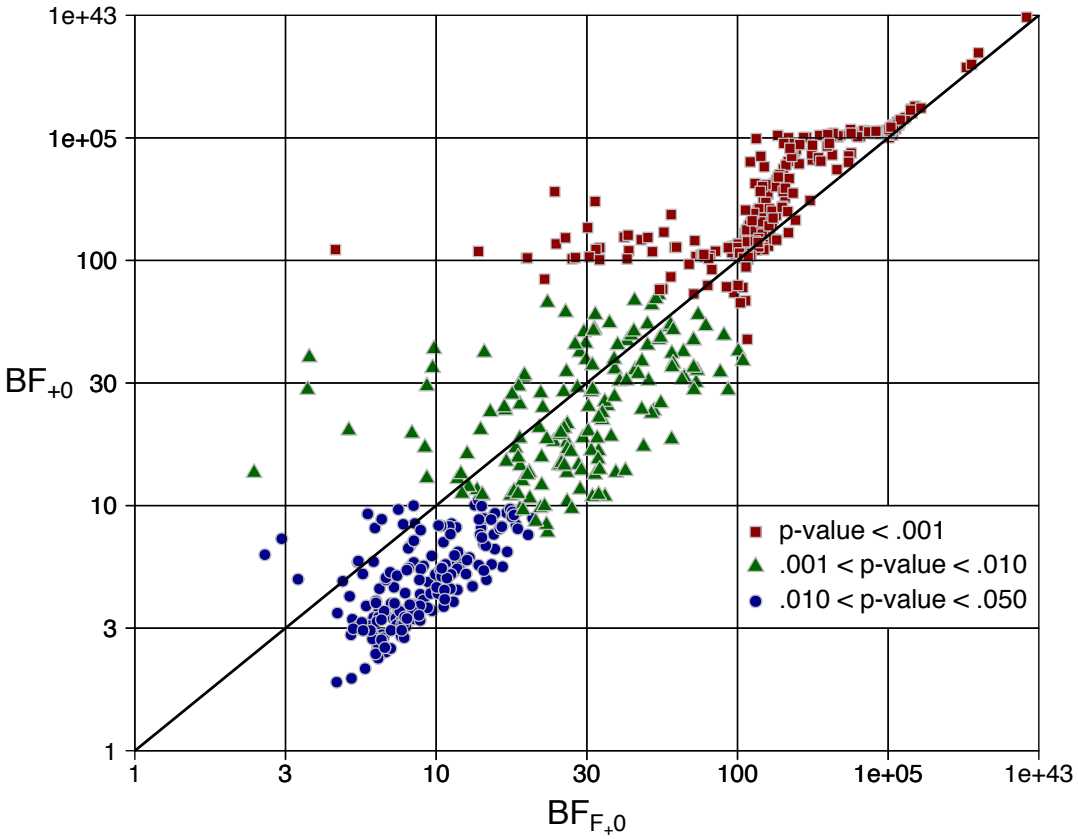

Figure 7: Comparison of the one-sided informed Bayes factor that specifies a zero trunctated $t(0.350, 0.102, 3)$ prior for δ under the alternative hypothesis and the one-sided default Bayes factor that specifies a zero truncated Cauchy$(0, 1/\sqrt{2})$ prior for δ under the alternative hypothesis for the 593 significant t-tests reported in Wetzels et al. (2011). The corresponding two-sided p-values are indicated by different shapes and colors: red squares correspond to p-values smaller than .001; green triangles correspond to p-values between .001 and .01; blue circles correspond to p-values between .01 and .05. Figure available at `https://tinyurl.com/m5fatwu` under CC license `https://creativecommons.org/licenses/by/2.0/`.

for Bayes factors that roughly correspond to t-tests with p-values in between .001 and .05, the informed Bayes factor generally provides more evidence for the alternative hypothesis than the default Bayes factor. For large Bayes factors that roughly correspond to p-values smaller than .001, this pattern is reversed.

This reversal can be explained by considering how well the informed alternative hypothesis has predicted the observed effect sizes. When considering the informed $t(0.350, 0.102, 3)$ prior, the central 95% prior credible interval ranges from about 0.025 to 0.675. Hence, when the observed effect sizes fall in the predicted range (i.e., roughly from 0.025 to 0.675) the informed alternative hypothesis will perform relatively well; when the observed effect sizes are larger or smaller the default alternative hypothesis will perform relatively well.

Figure 8 displays the same Bayes factors as Figure 7 but indicates the observed effect sizes by different shapes and colors. For all observed effect sizes that fall within the interval that is plausible under the informed alternative hypothesis, $\text{BF}_{F_{+}0}$ provides more evidence for an effect than BF_{+0}. Hence, this empirical investigation suggests that the informed generalization of Jeffreys's t-tests can lead to more diagnostic tests.

6 QUANTIFYING EVIDENCE FOR $\mathcal{H}_0$

A remaining question is how the informed Bayes factor compares to the default Bayes factor in case the null hypothesis is (approximately) true. To investigate this scenario we computed the maximum possible Bayes factor in favor of the null hypothesis as a function of the number of participants per group by fixing the t-value to zero. Figure 9 displays the results. When the number of participants per group is small, the default Bayes factor provides more evidence for $\mathcal{H}_0$ than does the informed Bayes factor. However, when the number of participants per group is larger than 82, this pattern is reversed.

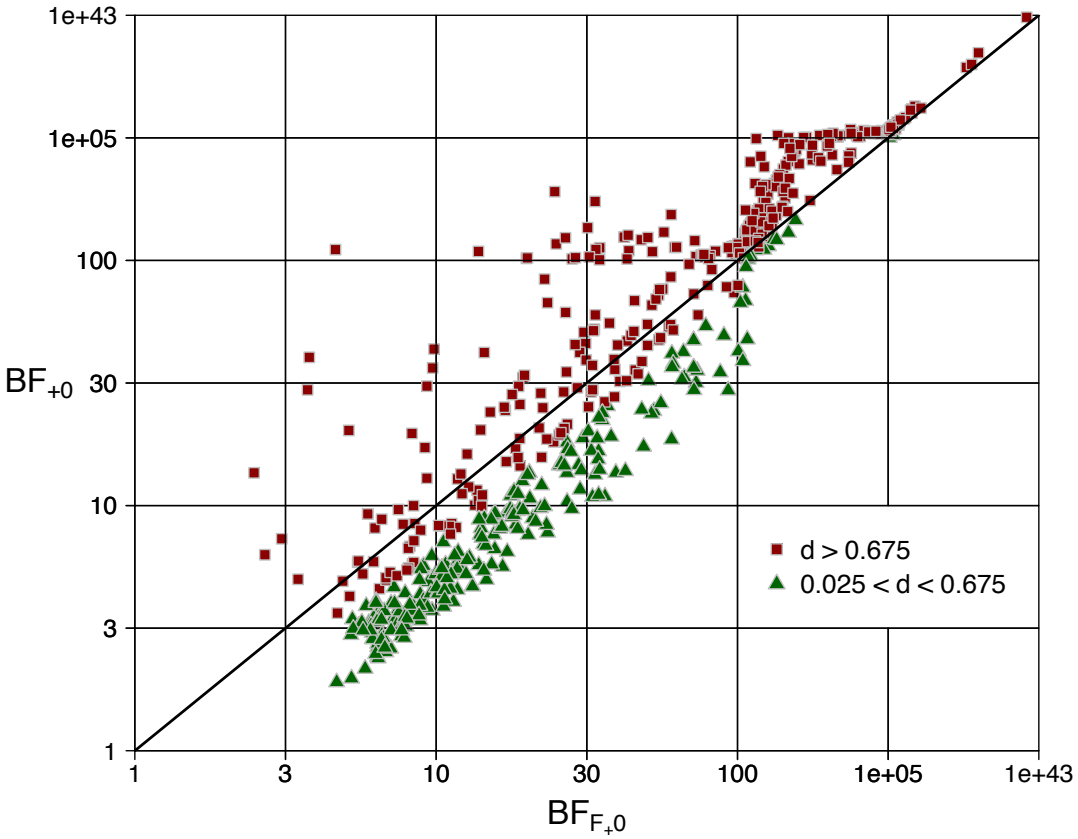

Figure 8: Comparison of the one-sided informed Bayes factor that specifies a zero truncated $t(0.350, 0.102, 3)$ prior for δ under the alternative hypothesis and the one-sided default Bayes factor that specifies a zero truncated Cauchy$(0, 1/\sqrt{2})$ prior for δ under the alternative hypothesis for the 593 significant t-tests reported in Wetzels et al. (2011). The corresponding observed effect sizes are indicated by different shapes and colors: red squares correspond to observed effect sizes larger than 0.675; green triangles correspond to observed effect sizes between 0.025 and 0.675. Figure available at `https://tinyurl.com/mo5g9y2` under CC license `https://creativecommons.org/licenses/by/2.0/`.

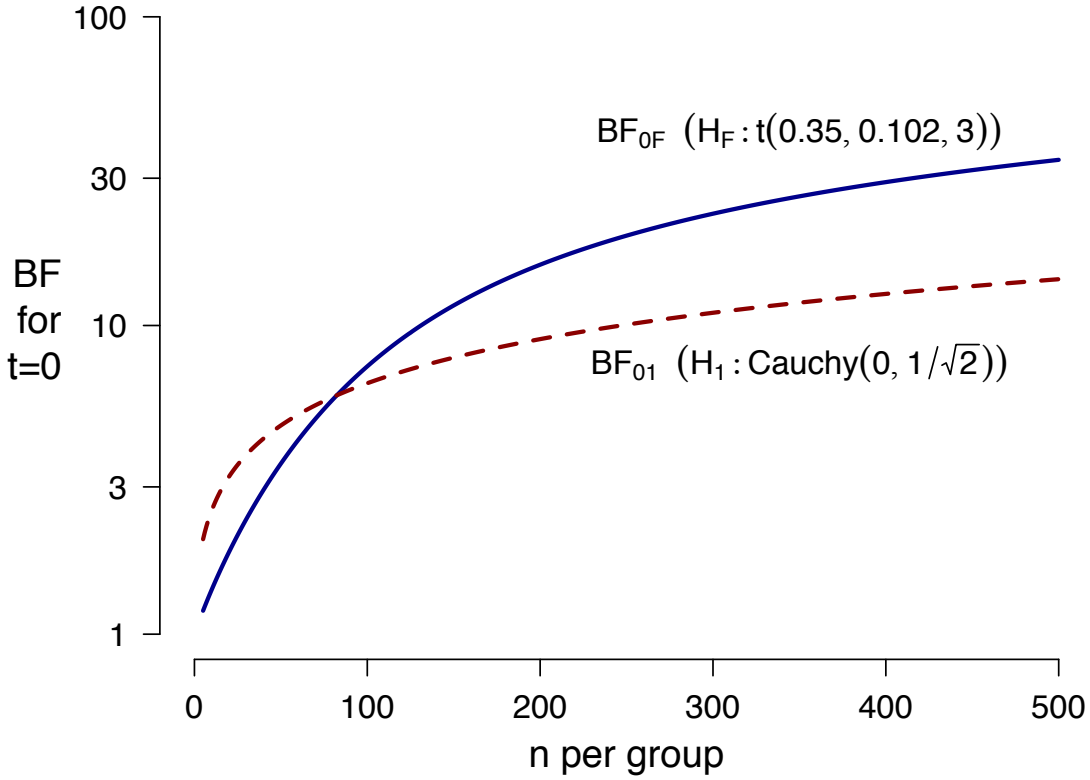

Figure 9: The maximum possible Bayes factor in favor of $\mathcal{H}_0$ as a function of the number of participants per group for unpaired t-tests. The dashed red line corresponds to the default Bayes factor that specifies a Cauchy$(0, 1/\sqrt{2})$ prior for δ under the alternative hypothesis; the solid blue line corresponds to the informed Bayes factor that specifies a $t(0.350, 0.102, 3)$ prior for δ under the alternative hypothesis. Figure available at `https://tinyurl.com/m92mvpx` under CC license `https://creativecommons.org/licenses/by/2.0/`.

An intuitive explanation for the reversal is provided by the Savage-Dickey density ratio representation of the Bayes factor (Dickey and Lientz, 1970): the Bayes factor in favor of $\mathcal{H}_0$ equals the ratio of posterior to prior density for δ under the alternative hypothesis evaluated at the test value $\delta = 0$. When n is small, the posterior for δ under the informed hypothesis is similar to the informed prior whereas the less restrictive default prior will be updated more strongly by the data. Hence, in case $\mathcal{H}_0$ is true, the ratio of posterior to prior density evaluated at zero will be larger for the default alternative hypothesis than for the informed alternative hypothesis. However, when n grows large, the data start to overwhelm the prior so that the posterior distributions become more similar. In that case, the ratio of posterior to prior density will be larger for the informed alternative hypothesis since its prior density at zero is smaller than that of the default prior. In the limit as $n \to \infty$, the posterior distributions will be identical; in this case, the difference of the log Bayes factor for the informed and the default prior is equal to the difference of the log of the informed prior evaluated at zero and the log of the default prior evaluated at zero (e.g., Gronau and Wagenmakers, 2017).

7 CONCLUDING COMMENTS

The comparison between two means is a quintessential inference problem. Harold Jeffreys believed that the only satisfactory solution lay in the application of a Bayes factor t-test (Rouder et al., 2009; Jeffreys, 1961; Ly et al., 2016b; Wetzels et al., 2009), the purpose of which is to quantify the relative predictive performance of two competing hypotheses: the null hypothesis $\mathcal{H}_0$ that represents an invariance or general law and the alternative hypothesis $\mathcal{H}_1$ that relaxes this restriction. However, in their current form, Jeffreys's tests do not allow researchers to incorporate expert knowledge in the specification of the prior

distribution for effect size under the alternative hypothesis, as all of Jeffreys' priors are centered on zero.

In this article, we have presented an informed extension of Jeffreys's t-tests that admits the specification of prior distributions centered away from zero. This flexibility may encourage the use of prior distributions that better represent the predictions from the hypothesis under test, producing more diagnostic outcomes for the same data. Specifically, our final example showed that when the observed effect size falls within the region of plausible values predicted by the informed prior, the informed Bayes factor provides more evidence for an effect than does the default Bayes factor. Furthermore, we illustrated that when $\mathcal{H}_0$ is true and the sample size is sufficiently large, the informed Bayes factor can also provide more evidence for the nullity of an effect than does the default Bayes factor.

The merits and debits of an informed prior specification will always remain a topic of debate. In our opinion, the prior distribution is an integral part of the model specification (e.g., Vanpaemel, 2010) and should be adjusted to include existing knowledge and impose meaningful constraints. The perceived dangers of the subjectivity that accompanies an informed analysis can be mitigated by preregistering the prior distribution (Chambers, 2013) and by conducting sensitivity analyses in which different prior choices are explored. For instance, when it comes to small-to-medium effects in experimental psychology, we believe the Oosterwijk prior is eminently plausible; different specifications are certainly possible, but they will need to respect the basic constraints (i.e., a mean around .35, most mass between .1 and .6) so that the results will be relatively insensitive to these choices.

We hope that empirical scientists will feel encouraged to pay more attention to the predictions of the hypotheses they seek to test (Rouder et al., 2016a,b). There is a price to be paid in order to apply an informed analysis –the method cannot be applied without thought– but the corresponding reward is a more diagnostic test.

References

Bateman, H., Erdélyi, A., Magnus, W., Oberhettinger, F., Tricomi, F. G., Bertin, D., Fulks, W. B., Harvey, A. R., Thomsen Jr, D. L., Weber, M. A., Whitney, E. L., and Stampfel, R. (1954). *Tables of integral transforms*, volume 1. McGraw-Hill.

Bayarri, M. J., Berger, J. O., Forte, A., and García-Donato, G. (2012). Criteria for Bayesian model choice with application to variable selection. *The Annals of Statistics*, 40:1550–1577.

Berger, J. O. and Berry, D. A. (1988). The relevance of stopping rules in statistical inference. In Gupta, S. S. and Berger, J. O., editors, *Statistical Decision Theory and Related Topics: Vol. 4*, pages 29–72. Springer Verlag, New York.

Chambers, C. D. (2013). Registered Reports: A new publishing initiative at Cortex. *Cortex*, 49:609–610.

Dickey, J. M. and Lientz, B. P. (1970). The weighted likelihood ratio, sharp hypotheses about chances, the order of a Markov chain. *The Annals of Mathematical Statistics*, 41:214–226.

Etz, A. and Wagenmakers, E.-J. (2017). J. B. S. Haldane's contribution to the Bayes factor hypothesis test. *In press for Statistical Science*.

Galton, F. (1907). Vox populi. *Nature*, 75:450–451.

Gronau, Q. F. and Wagenmakers, E.-J. (2017). Bayesian evidence accumulation in experimental mathematics: A case study of four irrational numbers. *In press for Experimental Mathematics*.

Haldane, J. B. S. (1932). A note on inverse probability. *Mathematical Proceedings of the Cambridge Philosophical Society*, 28:55–61.

Hoeting, J. A., Madigan, D., Raftery, A. E., and Volinsky, C. T. (1999). Bayesian model averaging: A tutorial. *Statistical Science*, 14:382–417.

JASP Team (2017). JASP (Version 0.8.1.1)[Computer software].

Jeffreys, H. (1942). On the significance tests for the introduction of new functions to represent measures. *Proceedings of the Royal Society (London) A: Mathematical, Physical and Engineering Sciences*, 180:256–268.

Jeffreys, H. (1948). *Theory of Probability*. Oxford University Press, Oxford, UK, 2 edition.

Jeffreys, H. (1961). *Theory of Probability*. Oxford University Press, Oxford, UK, 3 edition.

Kass, R. E. and Raftery, A. E. (1995). Bayes factors. *Journal of the American Statistical Association*, 90:773–795.

Liang, F., Paulo, R., Molina, G., Clyde, M. A., and Berger, J. O. (2008). Mixtures of g priors for Bayesian variable selection. *Journal of the American Statistical Association*, 103:410–423.

Lindley, D. V. (1972). *Bayesian Statistics, a Review*. SIAM, Philadelphia (PA).

Ly, A., Marsman, M., and Wagenmakers, E.-J. (2016a). Analytic posteriors for Pearson's correlation coefficient. *ArXiv*. https://arxiv.org/abs/1510.01188.

Ly, A., Verhagen, A. J., and Wagenmakers, E.-J. (2016b). Harold Jeffreys's default Bayes factor hypothesis tests: Explanation, extension, and application in psychology. *Journal of Mathematical Psychology*, 72:19–32.

Marsman, M., Schönbrodt, F. D., Morey, R. D., Yao, Y., Gelman, A., and Wagenmakers, E.-J. (2017). A Bayesian bird's eye view of 'Replications of important results in social psychology'. *Royal Society Open Science*, 4:160426.

Morey, R. D. and Rouder, J. N. (2015). BayesFactor 0.9.11-1. Comprehensive R Archive Network.

Morey, R. D. and Wagenmakers, E.-J. (2014). Simple relation between Bayesian order-restricted and point-null hypothesis tests. *Statistics and Probability Letters*, 92:121–124.

Morris, D. E., Oakley, J. E., and Crowe, J. A. (2014). A web-based tool for eliciting probability distributions from experts. *Environmental Modelling & Software*, 52:1–4.

Nuijten, M. B., Hartgerink, C. H., Assen, M. A., Epskamp, S., and Wicherts, J. M. (2016). The prevalence of statistical reporting errors in psychology (1985–2013). *Behavior research methods*, 48:1205–1226.

R Core Team (2016). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria.

Robert, C. P., Chopin, N., and Rousseau, J. (2009). Harold Jeffreys's Theory of Probability revisited. *Statistical Science*, 24:141–172.

Rouder, J. N. (2014). Optional stopping: No problem for Bayesians. *Psychonomic Bulletin & Review*, 21:301–308.

Rouder, J. N., Morey, R. D., Verhagen, J., Province, J. M., and Wagenmakers, E.-J. (2016a). Is there a free lunch in inference? *Topics in Cognitive Science*, 8:520–547.

Rouder, J. N., Morey, R. D., and Wagenmakers, E.-J. (2016b). The interplay between subjectivity, statistical practice, and psychological science. *Collabra*, 2:1–12.

Rouder, J. N., Speckman, P. L., Sun, D., Morey, R. D., and Iverson, G. (2009). Bayesian t tests for accepting and rejecting the null hypothesis. *Psychonomic Bulletin & Review*, 16:225–237.

Savage, L. J. (1961). The foundations of statistics reconsidered. In Neyman, J., editor, *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability, Vol. 1*, pages 575–586. University of California Press, Berkely, CA.

Steegen, S., Dewitte, L., Tuerlinckx, F., and Vanpaemel, W. (2014). Measuring the crowd within again: A pre-registered replication study. *Frontiers in Psychology*, 5:786.

Strack, F., Martin, L. L., and Stepper, S. (1988). Inhibiting and facilitating conditions of the human smile: A nonobtrusive test of the facial feedback hypothesis. *Journal of Personality and Social Psychology*, 54:768–777.

Surowiecki, J. (2004). *The Wisdom of Crowds: Why the Many Are Smarter Than the Few and How Collective Wisdom Shapes Business, Economies, Societies and Nations*. Doubleday, New York.

Vanpaemel, W. (2010). Prior sensitivity in theory testing: An apologia for the Bayes factor. *Journal of Mathematical Psychology*, 54:491–498.

Verhagen, A. J. and Wagenmakers, E.-J. (2014). Bayesian tests to quantify the result of a replication attempt. *Journal of Experimental Psychology: General*, 143:1457–1475.

Vul, E. and Pashler, H. (2008). Measuring the crowd within probabilistic representations within individuals. *Psychological Science*, 19:645–647.

Wagenmakers, E.-J., Beek, T., Dijkhoff, L., Gronau, Q. F., Acosta, A., Adams, R., ..., and Zwaan, R. A. (2016). Registered Replication Report: Strack, Martin, & Stepper (1988). *Perspectives on Psychological Science*, 11:917–928.

Wetzels, R., Matzke, D., Lee, M. D., Rouder, J. N., Iverson, G. J., and Wagenmakers, E.-J. (2011). Statistical evidence in experimental psychology: An empirical comparison using 855 t tests. *Perspectives on Psychological Science*, 6:291–298.

Wetzels, R., Raaijmakers, J. G. W., Jakab, E., and Wagenmakers, E.-J. (2009). How to quantify support for and against the null hypothesis: A flexible WinBUGS implementation of a default Bayesian t test. *Psychonomic Bulletin & Review*, 16:752–760.

Wrinch, D. and Jeffreys, H. (1921). On certain fundamental principles of scientific inquiry. *Philosophical Magazine*, 42:369–390.

8 Appendix A: Derivations One-Sample and Paired Samples T-Test

8.1 The Null model $\mathcal{H}_0$

8.1.1 Likelihood

We assume that we observe n observations which are drawn from a normal distribution parametrized in terms of the effect size $\delta = \frac{\mu}{\sigma}$, that is, $y_i \sim \mathcal{N}(\sigma\delta, \sigma^2) \, \forall i \in \{1, 2, \ldots, n\}$. Note that this includes the paired samples t-test case, since a paired samples t-test is equivalent to a one-sample t-test on the difference scores. The likelihood is given by:

$$p(y \mid \delta, \sigma^2) = (2\pi)^{-\frac{n}{2}} (\sigma^2)^{-\frac{n}{2}} \exp\left\{ -\frac{(n-1)s^2 + n(\bar{y} - \sigma\delta)^2}{2} \frac{1}{\sigma^2} \right\}. \tag{18}$$

For $\mathcal{H}_0$, we set $\delta = 0$ which yields:

$$p(y \mid \delta = 0, \sigma^2) = (2\pi)^{-\frac{n}{2}} (\sigma^2)^{-\frac{n}{2}} \exp\left\{ -\frac{(n-1)s^2 + n\bar{y}^2}{2} \frac{1}{\sigma^2} \right\}. \tag{19}$$

8.1.2 Prior Specification

Next we need to specify our prior distribution. There is only one parameter, σ^2, for which we will use Jeffreys's prior, that is, $p(\sigma^2) \propto \sigma^{-2}$.

8.1.3 Marginal Likelihood

In order to obtain the marginal likelihood, we need to integrate out the model parameter as follows:

$$\begin{aligned} p(y \mid \mathcal{H}_0) &= \int_0^\infty p(y \mid \delta = 0, \sigma^2)\, p(\sigma^2)\, \mathrm{d}\sigma^2 \\ &= (2\pi)^{-\frac{n}{2}} \int_0^\infty (\sigma^2)^{-\frac{n}{2}-1} \exp\left\{ -\frac{(n-1)s^2 + n\bar{y}^2}{2} \frac{1}{\sigma^2} \right\} \mathrm{d}\sigma^2. \end{aligned} \tag{20}$$

To solve this integral, we use the fact that the integrand is in the form of an unnormalized inverse-gamma distribution so that we know that the integral is equal to the normalizing constant of that inverse-gamma distribution, which yields:

$$p(y \mid \mathcal{H}_0) = (2\pi)^{-\frac{n}{2}} \Gamma\left(\frac{n}{2}\right) 2^{\frac{n}{2}} \left((n-1)s^2\right)^{-\frac{n}{2}} \left(1 + \frac{t^2}{n-1}\right)^{-\frac{n}{2}}. \tag{21}$$

8.2 $\mathcal{H}_1$: Shifted and Scaled t Distribution

8.2.1 Likelihood

It is convenient to rewrite the likelihood in such a way that we obtain the product of two terms where one term is equal to the pdf of a normal distribution and the other term is

independent of δ:

$$p(y \mid \delta, \sigma^2) = (2\pi\sigma^2)^{-\frac{n}{2}+\frac{1}{2}} n^{-\frac{1}{2}} \exp\left\{ -\frac{(n-1)s^2}{2\sigma^2} \right\} \times \left(2\pi\frac{\sigma^2}{n}\right)^{-\frac{1}{2}} \exp\left\{ -\frac{1}{2\frac{\sigma^2}{n}}(\bar{y} - \delta\sigma)^2 \right\}. \tag{22}$$

8.2.2 Prior Specification

For σ^2, we will use the same prior as in $\mathcal{H}_0$, that is, $p(\sigma^2) \propto \sigma^{-2}$. For the standardized effect size δ, we will use a shifted and scaled t distribution, that is $\delta \sim t(\mu_\delta, r, \kappa)$, where μ_δ corresponds to the location parameter, r to the scale parameter (note that we use r for consistency with Rouder et al., 2009; when the number of degrees of freedom is set to one, r corresponds to the Cauchy prior width), and κ to the degrees of freedom. For computing the marginal likelihood, it will be important to remember that we can write a t distribution as a scale mixture of normals (i.e., $\mathcal{N}(\mu_\delta, g)$), where we mix with respect to an inverse-gamma distribution of the form $g \sim IG\left(\frac{\kappa}{2}, r^2\frac{\kappa}{2}\right)$. This means that when we integrate over the auxiliary variable g, we obtain a $t(\mu_\delta, r, \kappa)$ distribution:

$$\begin{aligned} t(\delta; \mu_\delta, r, \kappa) &= \int_0^\infty \mathcal{N}(\delta; \mu_\delta, g)\, IG\left(g; \frac{\kappa}{2}, r^2\frac{\kappa}{2}\right) \mathrm{d}g \\ &= \frac{1}{r}\frac{\Gamma\left(\frac{\kappa+1}{2}\right)}{\sqrt{\pi\kappa}\,\Gamma(\frac{\kappa}{2})} \left[1 + \frac{\left(\frac{\delta-\mu_\delta}{r}\right)^2}{\kappa}\right]^{-\frac{\kappa+1}{2}}. \end{aligned} \tag{23}$$

8.2.3 Marginal Likelihood

In order to obtain the marginal likelihood, we need to integrate out the model parameters as follows:

$$\begin{aligned} p(y \mid \mathcal{H}_1) &= \int_0^\infty \int_0^\infty \int_{-\infty}^\infty p(y \mid \delta, \sigma^2)\, p(\delta \mid g)\, p(g)\, p(\sigma^2)\, \mathrm{d}\delta \mathrm{d}\sigma^2 \mathrm{d}g \\ &= \int_0^\infty p(g) \int_0^\infty n^{-\frac{1}{2}} (2\pi\sigma^2)^{-\frac{n}{2}+\frac{1}{2}} \exp\left\{ -\frac{(n-1)s^2}{2\sigma^2} \right\} \frac{1}{\sigma^2} \\ &\quad \times \int_{-\infty}^\infty \left(2\pi\frac{\sigma^2}{n}\right)^{-\frac{1}{2}} \exp\left\{ -\frac{1}{2\frac{\sigma^2}{n}} (\bar{y} - \delta\sigma)^2 \right\} \\ &\quad \times \left(2\pi g\right)^{-\frac{1}{2}} \exp\left\{ -\frac{1}{2g} (\delta - \mu_\delta)^2 \right\} \mathrm{d}\delta \mathrm{d}\sigma^2 \mathrm{d}g. \end{aligned} \tag{24}$$

We will first take a closer look at the integral with respect to δ. The first step is to make the following change of variable: $\mu = \delta\sigma$ with $\mathrm{d}\delta = \sigma^{-1}\mathrm{d}\mu$. We obtain:

$$\int_{-\infty}^\infty \left(2\pi\frac{\sigma^2}{n}\right)^{-\frac{1}{2}} \exp\left\{ -\frac{1}{2\frac{\sigma^2}{n}} (\bar{y} - \mu)^2 \right\} \left(2\pi g\sigma^2\right)^{-\frac{1}{2}} \exp\left\{ -\frac{1}{2g\sigma^2} \left(\mu - \sigma\mu_\delta\right)^2 \right\} \mathrm{d}\mu. \tag{25}$$

It is well known that this integral results in yet another normal distribution which is given by:

$$\left(2\pi\sigma^2 \left(\frac{1}{n} + g\right)\right)^{-\frac{1}{2}} \exp\left\{ -\frac{(\bar{y} - \sigma\mu_\delta)^2}{2\sigma^2 \left(\frac{1}{n} + g\right)} \right\}. \tag{26}$$

Next, we take a look at the integral with respect to σ^2. By using the just obtained result and rearranging terms, this integral reads as follows:

$$\begin{aligned} &(2\pi)^{-\frac{n-1}{2}} \left(2\pi\left(1 + ng\right)\right)^{-\frac{1}{2}} \exp\left\{ -\frac{\mu_\delta^2}{2\left(\frac{1}{n} + g\right)} \right\} \\ &\times \int_0^\infty \left(\sigma^2\right)^{-\frac{n}{2}-1} \exp\left\{ -\left[\frac{(n-1)s^2}{2} + \frac{\bar{y}^2}{2\left(\frac{1}{n} + g\right)} \right] \frac{1}{\sigma^2} + \frac{\bar{y}\mu_\delta}{\left(\frac{1}{n} + g\right)} \frac{1}{\sigma} \right\} \mathrm{d}\sigma^2. \end{aligned} \tag{27}$$

To evaluate this integral, we use a lemma presented in Ly et al. (2016a) which has been distilled from the Bateman project (Bateman et al., 1954).[1] Our final expression for the marginal likelihood of $\mathcal{H}_1$ is then given by:

$$\begin{aligned} p(y \mid \mathcal{H}_1) = \int_0^\infty (2\pi)^{-\frac{n-1}{2}} \left(2\pi\left(1+ng\right)\right)^{-\frac{1}{2}} \exp\left\{ -\frac{\mu_\delta^2}{2\left(\frac{1}{n}+g\right)}\right\} 2^{\frac{n}{2}} \left((n-1)s^2\right)^{-\frac{n}{2}} \\ \times \left(1+\frac{t^2}{(n-1)(1+ng)}\right)^{-\frac{n}{2}} [A+B]\, p(g) \mathrm{d}g, \end{aligned} \tag{28}$$

where $p(g)$ corresponds to the density of an inverse-gamma distribution of the form

$$p(g) = \frac{\left(r^2\frac{\kappa}{2}\right)^{\frac{\kappa}{2}}}{\Gamma(\frac{\kappa}{2})} g^{-\frac{\kappa}{2}-1} \exp\left(-\frac{r^2\kappa}{2g}\right), \tag{29}$$

and

$$A = \Gamma\left(\frac{n}{2}\right) {}_1F_1\left(\frac{n}{2}; \frac{1}{2}; \frac{\mu_\delta^2 t^2}{2\left(\frac{1}{n}+g\right)\left[(1+ng)\,(n-1)+t^2\right]}\right), \tag{30}$$

$$\begin{aligned} B = \frac{\mu_\delta t}{\sqrt{\frac{1}{2}\left(\frac{1}{n}+g\right)\left[(1+ng)\,(n-1)+t^2\right]}} \Gamma\left(\frac{n+1}{2}\right) \\ \times {}_1F_1\left(\frac{n+1}{2}; \frac{3}{2}; \frac{\mu_\delta^2 t^2}{2\left(\frac{1}{n}+g\right)\left[(1+ng)\,(n-1)+t^2\right]}\right), \end{aligned} \tag{31}$$

where ${}_1F_1$ corresponds to the confluent hypergeometric function.

8.3 Bayes Factor for a t Prior on δ

The Bayes factor BF_{10} is given by

$$\mathrm{BF}_{10} = \frac{p(y \mid \mathcal{H}_1)}{p(y \mid \mathcal{H}_0)}. \tag{32}$$

[1] Note that there are conditions that need to be fulfilled in order for this lemma to hold (specifically, certain components of the integrand need to be positive). In practical applications, those are almost always fulfilled (i.e., as long as there is more than one observation and the sample mean and sample variance are not both equal to zero, the lemma can be applied).

We notice that $2^{\frac{n}{2}}\left((n-1)s^2\right)^{-\frac{n}{2}}$ appears both in the marginal likelihood for $\mathcal{H}_1$ and $\mathcal{H}_0$ and thus cancels. The resulting Bayes factor is:

$$\text{BF}_{10} = \frac{\int_0^\infty (1+ng)^{-\frac{1}{2}} \exp\left\{-\frac{\mu_\delta^2}{2\left(\frac{1}{n}+g\right)}\right\}\left(1+\frac{t^2}{(n-1)(1+ng)}\right)^{-\frac{n}{2}} [A+B]\, p(g)\mathrm{d}g}{\Gamma\left(\frac{n}{2}\right)\left(1+\frac{t^2}{n-1}\right)^{-\frac{n}{2}}}. \tag{33}$$

Notice that the only information we need for computing this Bayes factor is the t-value and the sample size n. Furthermore, note that in the main text, we have expressed some of the n's with respect to the degrees of freedom $\nu = n-1$ for an easier extension of the expressions to the unpaired t-test case.

8.4 Bayes Factor for a Normal Prior on δ

The Bayes factor for a normal prior of the form $\mathcal{N}(\mu_\delta, g)$ for δ is obtained by simply removing the integral over g with respect to $p(g)$:

$$\text{BF}_{10} = \frac{(1+ng)^{-\frac{1}{2}} \exp\left\{-\frac{\mu_\delta^2}{2\left(\frac{1}{n}+g\right)}\right\}\left(1+\frac{t^2}{(n-1)(1+ng)}\right)^{-\frac{n}{2}} [A+B]}{\Gamma\left(\frac{n}{2}\right)\left(1+\frac{t^2}{n-1}\right)^{-\frac{n}{2}}}. \tag{34}$$

8.5 Marginal Posterior Distribution for δ

8.5.1 Shifted and Scaled t Prior

The marginal posterior distribution for δ under $\mathcal{H}_1$ is given by:

$$p(\delta \mid y) = \frac{\int_0^\infty p(y \mid \delta, \sigma^2)\, p(\sigma^2)\mathrm{d}\sigma^2\, p(\delta)}{p(y)}, \tag{35}$$

where $p(\delta)$ refers to the shifted and scaled t prior distribution for δ. Note that the denominator of this expression, $p(y)$, is simply the marginal likelihood of $\mathcal{H}_1$ which we have calculated previously. The integral in the numerator can be simplified as follows:

$$\begin{aligned}&\int_0^\infty p(y \mid \delta, \sigma^2)\, p(\sigma^2)\mathrm{d}\sigma^2 \\ =&(2\pi)^{-\frac{n}{2}} \exp\left\{-\frac{n}{2}\delta^2\right\} \int_0^\infty \left(\sigma^2\right)^{-\frac{n}{2}-1} \exp\left\{-\frac{1}{2}\left[(n-1)s^2 + n\bar{y}^2\right]\frac{1}{\sigma^2} + n\bar{y}\delta\frac{1}{\sigma}\right\}\mathrm{d}\sigma^2.\end{aligned} \tag{36}$$

This integral can be solved by using again the lemma presented in Ly et al. (2016a).[2] We obtain the following expression for the marginal posterior distribution for δ:

$$p(\delta \mid y) = \frac{\exp\left\{-\frac{n}{2}\delta^2\right\}\left(1+\frac{t^2}{n-1}\right)^{-\frac{n}{2}}[C+D]\,\frac{1}{r}\frac{\Gamma\left(\frac{\kappa+1}{2}\right)}{\sqrt{\pi\kappa}\,\Gamma\left(\frac{\kappa}{2}\right)}\left[1+\frac{\left(\frac{\delta-\mu_\delta}{r}\right)^2}{\kappa}\right]^{-\frac{\kappa+1}{2}}}{\int_0^\infty (1+ng)^{-\frac{1}{2}}\exp\left\{-\frac{\mu_\delta^2}{2\left(\frac{1}{n}+g\right)}\right\}\left(1+\frac{t^2}{(n-1)(1+ng)}\right)^{-\frac{n}{2}}[A+B]\,p(g)\mathrm{d}g}, \tag{37}$$

where

$$C = \Gamma\left(\frac{n}{2}\right){}_1F_1\left(\frac{n}{2};\frac{1}{2};\frac{nt^2\delta^2}{2(n-1+t^2)}\right), \tag{38}$$

$$D = t\delta\sqrt{\frac{2n}{n-1+t^2}}\Gamma\left(\frac{n+1}{2}\right){}_1F_1\left(\frac{n+1}{2};\frac{3}{2};\frac{nt^2\delta^2}{2(n-1+t^2)}\right). \tag{39}$$

Note that we only need to know the t-value and the sample size in order to calculate the posterior distribution.

Normal Prior

To obtain the marginal posterior distribution for δ for a normal prior, we simply need to replace the prior distribution for δ and remove the integral over g with respect to $p(g)$

[2] Again, note that the lemma can be applied as long as there is more than one observation and the sample mean and sample variance are not both equal to zero.

which yields:

$$p(\delta \mid y) = \frac{\exp\left\{-\frac{n}{2}\delta^2\right\}\left(1+\frac{t^2}{n-1}\right)^{-\frac{n}{2}}[C+D]\,(2\pi g)^{-\frac{1}{2}}\exp\left\{-\frac{1}{2g}(\delta-\mu_\delta)^2\right\}}{(1+ng)^{-\frac{1}{2}}\exp\left\{-\frac{\mu_\delta^2}{2\left(\frac{1}{n}+g\right)}\right\}\left(1+\frac{t^2}{(n-1)(1+ng)}\right)^{-\frac{n}{2}}[A+B]}. \tag{40}$$

Again, note that, in the main text, we have expressed some of the n's with respect to the degrees of freedom $\nu = n-1$ for an easier extension of the expressions to the unpaired t-test case.

9 Appendix B: Derivations Unpaired Samples T-Test

9.1 The Null model $\mathcal{H}_0$

9.1.1 Likelihood

We assume that we observe n_y observations which are drawn from a normal distribution parametrized in terms of the effect size $\delta = \frac{\alpha}{\sigma}$, that is, $y_i \sim \mathcal{N}(\mu - \frac{\sigma\delta}{2}, \sigma^2)\,\forall i \in \{1,2,\ldots,n_y\}$ in the first group, and n_x observations which are drawn from a normal distribution of the form $x_j \sim \mathcal{N}(\mu + \frac{\sigma\delta}{2}, \sigma^2)\,\forall j \in \{1,2,\ldots,n_x\}$ in the second group. Let $s_y^2 = \frac{1}{n_y-1}\sum_{i=1}^{n_y}(y_i-\bar{y})^2$ and $s_x^2 = \frac{1}{n_x-1}\sum_{j=1}^{n_x}(x_j-\bar{x})^2$. The likelihood is given by:

$$\begin{aligned} p(y,x \mid \delta,\mu,\sigma^2) =& (2\pi)^{-\frac{n_y+n_x}{2}}(\sigma^2)^{-\frac{n_y+n_x}{2}} \\ &\times \exp\left\{-\frac{1}{2\sigma^2}\left[(n_y-1)s_y^2 + n_y\left(\bar{y}-\left(\mu-\frac{\sigma\delta}{2}\right)\right)^2\right.\right. \\ &\left.\left. + (n_x-1)s_x^2 + n_x\left(\bar{x}-\left(\mu+\frac{\sigma\delta}{2}\right)\right)^2\right]\right\}. \end{aligned} \tag{41}$$

For the null model, we set $\delta = 0$ which yields:

$$\begin{aligned} p(y, x \mid \delta = 0, \mu, \sigma^2) =& (2\pi)^{-\frac{n_y+n_x}{2}} (\sigma^2)^{-\frac{n_y+n_x}{2}} \\ & \times \exp\left\{ -\frac{1}{2\sigma^2}\Big[(n_y-1)s_y^2 + n_y\left(\bar{y}-\mu\right)^2 \right. \\ & \left. \qquad + (n_x-1)s_x^2 + n_x\left(\bar{x}-\mu\right)^2\Big]\right\}. \end{aligned} \tag{42}$$

9.1.2 Prior Specification

Next we need to specify our prior distribution. For the parameters that are common to $\mathcal{H}_0$ and $\mathcal{H}_1$, we will use Jeffreys's prior, that is, $p(\mu, \sigma^2) \propto \sigma^{-2}$.

9.1.3 Marginal Likelihood

In order to obtain the marginal likelihood, we need to integrate out the model parameters as follows:

$$\begin{aligned} p(y, x \mid \mathcal{H}_0) &= \int_0^\infty \int_{-\infty}^\infty p(y, x \mid \delta = 0, \mu, \sigma^2)\, p(\mu, \sigma^2)\, \mathrm{d}\mu \mathrm{d}\sigma^2 \\ &= (2\pi)^{-\frac{n_y+n_x}{2}} \int_0^\infty (\sigma^2)^{-\frac{n_y+n_x}{2}-1} \exp\left\{-\frac{1}{2\sigma^2}\left[(n_y-1)s_y^2 + (n_x-1)s_x^2\right]\right\} \\ &\quad \times \int_{-\infty}^\infty \exp\left\{ -\frac{1}{2\sigma^2}\left[n_y\left(\bar{y}-\mu\right)^2 + n_x\left(\bar{x}-\mu\right)^2\right]\right\} \mathrm{d}\mu \mathrm{d}\sigma^2. \end{aligned} \tag{43}$$

We will first concentrate on the innermost integral. In order to solve the integral with respect to μ, we complete the square in μ and then recognize that we obtain an unnormalized normal distribution for which we know the normalizing constant which yields:

$$\exp\left\{ -\frac{1}{2\sigma^2}\left[n_y\bar{y}^2 + n_x\bar{x}^2 - \frac{\left(n_y\bar{y} + n_x\bar{x}\right)^2}{n_y+n_x}\right]\right\} \left(2\pi\frac{\sigma^2}{n_y+n_x}\right)^{\frac{1}{2}}. \tag{44}$$

Next, we solve the integral with respect to σ^2 which is accomplished by rearranging terms so that we obtain an unnormalized inverse-gamma density for which we know the normalizing

constant and we obtain:

$$\begin{aligned} p(y, x \mid \mathcal{H}_0) =& (2\pi)^{-\frac{n_y+n_x}{2}} \left(2\pi \frac{1}{n_y+n_x}\right)^{\frac{1}{2}} \Gamma\left(\frac{n_y+n_x-1}{2}\right) 2^{\frac{n_y+n_x-1}{2}} \\ & \times \left[(n_y-1)s_y^2 + (n_x-1)s_x^2\right]^{-\frac{n_y+n_x-1}{2}} \left[1 + \frac{t^2}{n_y+n_x-2}\right]^{-\frac{n_y+n_x-1}{2}}. \end{aligned} \tag{45}$$

9.2 $\mathcal{H}_1$: Shifted and Scaled t Distribution

9.2.1 Likelihood

To make the following derivation easier, we will first manipulate the likelihood so that we obtain a convenient expression:

$$\begin{aligned} p(y, x \mid \delta, \mu, \sigma^2) =& (2\pi)^{-\frac{n_y+n_x}{2}} (\sigma^2)^{-\frac{n_y+n_x}{2}} \\ & \times \exp\left\{ -\frac{1}{2\sigma^2}\left[\sum_{i=1}^{n_y} y_i^2 - 2\left(\mu - \frac{\sigma\delta}{2}\right) n_y \bar{y} + n_y \left(\mu - \frac{\sigma\delta}{2}\right)^2 \right.\right. \\ & \left.\left. + \sum_{j=1}^{n_x} x_j^2 - 2\left(\mu + \frac{\sigma\delta}{2}\right) n_x \bar{x} + n_x \left(\mu + \frac{\sigma\delta}{2}\right)^2 \right] \right\}. \end{aligned} \tag{46}$$

Let $n = n_y + n_x$ be the combined sample size and $\xi = \frac{\sigma\delta}{2}$. Furthermore, let $\bar{z} = \frac{\sum_{i=1}^{n_y} y_i + \sum_{j=1}^{n_x} x_j}{n}$ be the overall mean and $s_z^2 = \frac{\sum_{i=1}^{n_y}(y_i-\bar{z})^2 + \sum_{j=1}^{n_x}(x_j-\bar{z})^2}{n}$ be the overall average sum of squares. Note that $\sum_{i=1}^{n_y} y_i^2 + \sum_{j=1}^{n_x} x_j^2 = n s_z^2 + n\bar{z}^2$. Let $q_y = \frac{n_y}{n}$ and $q_x = \frac{n_x}{n}$ be the proportion of the overall samples size in group y and group x, respectively (hence, $q_y + q_x = 1$). Then, we can rewrite the likelihood as

$$p(y, x \mid \delta, \mu, \sigma^2) = (2\pi)^{-\frac{n}{2}} (\sigma^2)^{-\frac{n}{2}} \exp\left\{ -\frac{n}{2\sigma^2}\left[s_z^2 + 4[q_x(\bar{z}-\bar{x})]\xi + 4q_y q_x \xi^2 + (\mu - [\bar{z} + (q_y - q_x)\xi])^2 \right] \right\}. \tag{47}$$

9.2.2 Prior Specification

For μ and σ^2, we will use the same prior as in $\mathcal{H}_0$, that is, $p(\mu, \sigma^2) \propto \sigma^{-2}$. For the standardized effect size δ, we will use a shifted and scaled t distribution, that is $\delta \sim t(\mu_\delta, r, \kappa)$, where μ_δ corresponds to the location parameter, r to the scale parameter, and κ to the degrees of freedom parameter. As in the one-sample case, we make use of the fact that we can write a t distribution as a scale mixture of normals (i.e., $\mathcal{N}(\mu_\delta, g)$), where we mix with respect to an inverse-gamma distribution of the form $g \sim IG\left(\frac{\kappa}{2}, r^2\frac{\kappa}{2}\right)$.

9.2.3 Marginal Likelihood

In order to obtain the marginal likelihood, we need to integrate out the model parameters as follows:

$$\begin{aligned} p(y, x \mid \mathcal{H}_1) = &\int_0^\infty (2\pi)^{-\frac{n}{2}} p(g) \int_0^\infty (\sigma^2)^{-\frac{n}{2}-1} \\ &\times \int_{-\infty}^\infty \exp\left\{ -\frac{n}{2\sigma^2}\left[s_z^2 + 4[q_x(\bar{z} - \bar{x})]\xi + 4q_y q_x \xi^2\right]\right\} \\ &\times (2\pi g)^{-\frac{1}{2}} \exp\left\{-\frac{1}{2g}\left(\delta - \mu_\delta\right)^2\right\} \\ &\times \int_{-\infty}^\infty \exp\left\{ -\frac{n}{2\sigma^2}\left[(\mu - [\bar{z} + (q_y - q_x)\xi])^2\right]\right\} \mathrm{d}\mu \mathrm{d}\delta \mathrm{d}\sigma^2 \mathrm{d}g. \end{aligned} \tag{48}$$

We recognize that the innermost integral with respect to μ has the form of an unnormalized normal distribution. Hence, we obtain:

$$\begin{aligned} p(y, x \mid \mathcal{H}_1) = &\int_0^\infty (2\pi)^{\frac{1}{2}} n^{-\frac{1}{2}} (2\pi)^{-\frac{n}{2}} p(g) \int_0^\infty (\sigma^2)^{-\frac{n-1}{2}-1} \\ &\times \int_{-\infty}^\infty \exp\left\{ -\frac{n}{2\sigma^2}\left[s_z^2 + 4[q_x(\bar{z} - \bar{x})]\xi + 4q_y q_x \xi^2\right]\right\} \\ &\times (2\pi g)^{-\frac{1}{2}} \exp\left\{-\frac{1}{2g}\left(\delta - \mu_\delta\right)^2\right\} \mathrm{d}\delta \mathrm{d}\sigma^2 \mathrm{d}g. \end{aligned} \tag{49}$$

The next step is to integrate out δ. In order to achieve this, we first substitute ξ with $\frac{\sigma\delta}{2}$, then, we do a change-of-variable of the form $\alpha = \sigma\delta$ with $(\sigma^2)^{-\frac{1}{2}}\,\mathrm{d}\alpha = \mathrm{d}\delta$, and finally, we complete the square in α, recognize the result as an unnormalized normal distribution for which we know the normalizing constant, and rearrange, which yields:

$$\begin{aligned}&\int_0^\infty (2\pi)^{-\frac{n-1}{2}} n^{-\frac{1}{2}} g^{-\frac{1}{2}} \left(nq_yq_x + \frac{1}{g}\right)^{-\frac{1}{2}} \exp\left\{-\frac{\mu_\delta^2}{2\left(\frac{n_y+n_x}{n_yn_x}+g\right)}\right\} p(g)\\ &\times \int_0^\infty (\sigma^2)^{-\frac{n-1}{2}-1} \exp\left\{-\frac{n}{2}\left[s_z^2 - \frac{nq_x^2(\bar{z}-\bar{x})^2}{nq_yq_x+\frac{1}{g}}\right]\frac{1}{\sigma^2} - \frac{\mu_\delta nq_x(\bar{z}-\bar{x})}{gnq_yq_x+1}\frac{1}{\sigma}\right\}\mathrm{d}\sigma^2\mathrm{d}g.\end{aligned} \tag{50}$$

To evaluate the integral with respect to σ^2, we again use the lemma presented in Ly et al. (2016a).[3] The marginal likelihood can then be written as

$$\begin{aligned}p(y,x \mid \mathcal{H}_1) = &\int_0^\infty n^{-\frac{1}{2}}(2\pi)^{-\frac{n-1}{2}} \left(1+\frac{n_yn_x}{n_y+n_x}g\right)^{-\frac{1}{2}} \exp\left\{-\frac{\mu_\delta^2}{2\left(\frac{n_y+n_x}{n_yn_x}+g\right)}\right\} p(g)\\ &\times 2^{\frac{n-1}{2}}\left[(n_y-1)s_y^2+(n_x-1)s_x^2\right]^{-\frac{n-1}{2}} \left[1+\frac{t^2}{(n_y+n_x-2)\left(1+g\frac{n_yn_x}{n_y+n_x}\right)}\right]^{-\frac{n-1}{2}}\\ &\times [A+B]\,\mathrm{d}g,\end{aligned} \tag{51}$$

where

$$A = \Gamma\left(\frac{n-1}{2}\right) {}_1F_1\left(\frac{n-1}{2}; \frac{1}{2}; \frac{\mu_\delta^2 t^2}{2\left(\frac{n_y+n_x}{n_yn_x}+g\right)\left[(n_y+n_x-2)\left(1+\frac{n_yn_x}{n_y+n_x}g\right)+t^2\right]}\right), \tag{52}$$

[3]The lemma can be applied for all practically interesting cases (i.e., there is more than one observation in each group and the sample variances are non-zero).

$$B = \frac{\mu_\delta t}{\sqrt{\frac{1}{2}\left(\frac{n_y+n_x}{n_y n_x}+g\right)\left[(n_y+n_x-2)\left(1+\frac{n_y n_x}{n_y+n_x}g\right)+t^2\right]}}\Gamma\left(\frac{n}{2}\right) \\ \times {}_1F_1\left(\frac{n}{2};\frac{3}{2};\frac{\mu_\delta^2 t^2}{2\left(\frac{n_y+n_x}{n_y n_x}+g\right)\left[(n_y+n_x-2)\left(1+\frac{n_y n_x}{n_y+n_x}g\right)+t^2\right]}\right). \tag{53}$$

9.3 Bayes Factor for a t Prior on δ

The Bayes factor BF_{10} is given by

$$\text{BF}_{10} = \frac{p(y,x \mid \mathcal{H}_1)}{p(y,x \mid \mathcal{H}_0)}. \tag{54}$$

We notice that $(2\pi)^{-\frac{n-1}{2}} n^{-\frac{1}{2}} 2^{\frac{n-1}{2}} \left[(n_y-1)s_y^2+(n_x-1)s_x^2\right]^{-\frac{n-1}{2}}$ appears both in the marginal likelihood for $\mathcal{H}_0$ and $\mathcal{H}_1$ and thus cancels. Hence, the Bayes factor is given by:

$$\text{BF}_{10} = \\ \frac{\int_0^\infty \left(1+\frac{n_y n_x}{n_y+n_x}g\right)^{-\frac{1}{2}} \exp\left\{-\frac{\mu_\delta^2}{2\left(\frac{n_y+n_x}{n_y n_x}+g\right)}\right\}\left[1+\frac{t^2}{(n_y+n_x-2)\left(1+g\frac{n_y n_x}{n_y+n_x}\right)}\right]^{-\frac{n_y+n_x-1}{2}}[A+B]\,p(g)\mathrm{d}g}{\Gamma\left(\frac{n_y+n_x-1}{2}\right)\left[1+\frac{t^2}{n_y+n_x-2}\right]^{-\frac{n_y+n_x-1}{2}}}. \tag{55}$$

Note that to calculate the Bayes factor, we only need to know the t-value and the sample sizes.

9.4 Bayes Factor for a Normal Prior on δ

The Bayes factor for a normal prior of the form $\mathcal{N}(\mu_\delta, g)$ for δ is obtained by simply removing the integral over g with respect to $p(g)$:

$$\mathrm{BF}_{10} = \frac{\left(1 + \frac{n_y n_x}{n_y+n_x} g\right)^{-\frac{1}{2}} \exp\left\{ - \frac{\mu_\delta^2}{2\left(\frac{n_y+n_x}{n_y n_x}+g\right)} \right\} \left[1 + \frac{t^2}{(n_y+n_x-2)\left(1+g\frac{n_y n_x}{n_y+n_x}\right)}\right]^{-\frac{n_y+n_x-1}{2}} [A+B]}{\Gamma\left(\frac{n_y+n_x-1}{2}\right)\left[1+\frac{t^2}{n_y+n_x-2}\right]^{-\frac{n_y+n_x-1}{2}}}. \tag{56}$$

Note that, in the main text, we have expressed some of the terms that involve $n_y + n_x$ in terms of the degrees of freedom $\nu = n_y + n_x - 2$ and we have defined an effective sample size $n_{\text{eff}} = \frac{n_y n_x}{n_y+n_x}$ which allowed us to highlight similarities to the one-sample/paired samples results.

9.5 Marginal Posterior Distribution for δ

9.5.1 Shifted and Scaled t Prior

The marginal posterior distribution for δ under $\mathcal{H}_1$ is given by:

$$p(\delta \mid y, x) = \frac{\int_0^\infty \int_{-\infty}^\infty p(y, x \mid \delta, \mu, \sigma^2)\, p(\mu, \sigma^2) \mathrm{d}\mu \mathrm{d}\sigma^2\, p(\delta)}{p(y, x)}, \tag{57}$$

where $p(\delta)$ refers to the shifted and scaled t prior distribution for δ. Note that the denominator of this expression, $p(y, x)$, is simply the marginal likelihood of $\mathcal{H}_1$ which we have calculated previously. To evaluate the integrals in the numerator, we first notice that we have already shown how to integrate out μ which leaves us with the following integral with

respect to σ^2:

$$\begin{aligned}
&\int_0^\infty \int_{-\infty}^\infty p(y, x \mid \delta, \mu, \sigma^2)\, p(\mu, \sigma^2) \mathrm{d}\mu \mathrm{d}\sigma^2 \\
=&(2\pi)^{\frac{1}{2}} n^{-\frac{1}{2}} (2\pi)^{-\frac{n}{2}} \exp\left\{-\frac{n}{2} q_y q_x \delta^2\right\} \\
&\times \int_0^\infty (\sigma^2)^{-\frac{n-1}{2}-1} \exp\left\{ -\frac{n}{2} s_z^2 \frac{1}{\sigma^2} - q_x(\bar{z} - \bar{x})\delta \frac{1}{\sigma} \right\} \mathrm{d}\sigma^2.
\end{aligned} \tag{58}$$

To evaluate this integral, we use once more the lemma presented in Ly et al. (2016a).[4] Our final expression for the evaluated integral is

$$\begin{aligned}
&(n_y + n_x)^{-\frac{1}{2}} (2\pi)^{-\frac{n_y+n_x-1}{2}} \exp\left\{-\frac{n_y n_x}{2(n_y+n_x)} \delta^2\right\} 2^{\frac{n_y+n_x-1}{2}} \\
&\times \left[(n_y - 1)s_y^2 + (n_x - 1)s_x^2\right]^{-\frac{n_y+n_x-1}{2}} \left[1 + \frac{t^2}{n_y + n_x - 2}\right]^{-\frac{n_y+n_x-1}{2}} [C + D],
\end{aligned} \tag{59}$$

where

$$C = \Gamma\left(\frac{n_y + n_x - 1}{2}\right) {}_1F_1\left(\frac{n_y + n_x - 1}{2}; \frac{1}{2}; \frac{\frac{n_y n_x}{n_y+n_x} t^2 \delta^2}{2\left(n_y + n_x - 2 + t^2\right)}\right), \tag{60}$$

$$D = t\delta \sqrt{\frac{2\frac{n_y n_x}{n_y+n_x}}{n_y + n_x - 2 + t^2}} \Gamma\left(\frac{n_y + n_x}{2}\right) {}_1F_1\left(\frac{n_y + n_x}{2}; \frac{3}{2}; \frac{\frac{n_y n_x}{n_y+n_x} t^2 \delta^2}{2\left(n_y + n_x - 2 + t^2\right)}\right). \tag{61}$$

Using this expression and using the previously obtained result for the marginal likelihoood for $\mathcal{H}_1$, we obtain the following expression for the marginal posterior distribution for δ (note that we already canceled

$$(2\pi)^{-\frac{n_y+n_x-1}{2}} (n_y + n_x)^{-\frac{1}{2}} 2^{\frac{n_y+n_x-1}{2}} \left[(n_y - 1)s_y^2 + (n_x - 1)s_x^2\right]^{-\frac{n_y+n_x-1}{2}} \tag{62}$$

[4] Again, the lemma can be applied for all practically interesting cases (i.e., there is more than one observation in each group and the sample variances are non-zero).

which appeared both in the expression we have just obtained and the marginal likelihood for $\mathcal{H}_1$):

$$p(\delta \mid y, x) = \frac{\exp\left\{-\frac{n_y n_x}{2(n_y+n_x)}\delta^2\right\}\left[1+\frac{t^2}{n_y+n_x-2}\right]^{-\frac{n_y+n_x-1}{2}}[C+D]\,\frac{1}{r}\frac{\Gamma\left(\frac{\kappa+1}{2}\right)}{\sqrt{\pi\kappa}\,\Gamma\left(\frac{\kappa}{2}\right)}\left[1+\frac{\left(\frac{\delta-\mu_\delta}{r}\right)^2}{\kappa}\right]^{-\frac{\kappa+1}{2}}}{\int_0^\infty\left(1+\frac{n_y n_x}{n_y+n_x}g\right)^{-\frac{1}{2}}\exp\left\{-\frac{\mu_\delta^2}{2\left(\frac{n_y+n_x}{n_y n_x}+g\right)}\right\}\left[1+\frac{t^2}{(n_y+n_x-2)\left(1+g\frac{n_y n_x}{n_y+n_x}\right)}\right]^{-\frac{n_y+n_x-1}{2}}[A+B]\,p(g)\mathrm{d}g}. \tag{63}$$

Note that we only need to know the t-value and the sample sizes in order to calculate the posterior distribution.

9.5.2 Normal Prior

To obtain the marginal posterior distribution for δ for a normal prior, we simply need to replace the prior distribution for δ and remove the integral over g with respect to $p(g)$ which yields:

$$p(\delta \mid y, x) = \frac{\exp\left\{-\frac{n_y n_x}{2(n_y+n_x)}\delta^2\right\}\left[1+\frac{t^2}{n_y+n_x-2}\right]^{-\frac{n_y+n_x-1}{2}}[C+D]\,(2\pi g)^{-\frac{1}{2}}\exp\left\{-\frac{1}{2g}(\delta-\mu_\delta)^2\right\}}{\left(1+\frac{n_y n_x}{n_y+n_x}g\right)^{-\frac{1}{2}}\exp\left\{-\frac{\mu_\delta^2}{2\left(\frac{n_y+n_x}{n_y n_x}+g\right)}\right\}\left[1+\frac{t^2}{(n_y+n_x-2)\left(1+g\frac{n_y n_x}{n_y+n_x}\right)}\right]^{-\frac{n_y+n_x-1}{2}}[A+B]}, \tag{64}$$

Again, note that, in the main text, we have expressed some of the terms that involve n_y+n_x in terms of the degrees of freedom $\nu = n_y+n_x-2$ and we have defined an effective sample size $n_{\text{eff}} = \frac{n_y n_x}{n_y+n_x}$ which allowed us to highlight similarities to the one-sample/paired samples results.

10 Appendix C: Prior Elicitation

Here we describe in more detail how we elicited the prior distribution from Dr. Suzanne Oosterwijk. We applied the "Roulette" method which presents the expert with a grid of equally sized bins that cover the plausible range of possible effect size values. The expert then has to allocate chips to the bins to construct the prior distribution. The MATCH Uncertainty Elicitation Tool provides an immediate feedback by showing a fitted parametric distribution that matches the distribution specified via the chips (in our case, a t distribution fit), and it also provides the user with the 33%-tile and 66%-tile values that correspond to the specified distribution. Figure 10 displays a screenshot of the Roulette method with feedback using the MATCH Uncertainty Elicitation Tool. The fitted t distribution that is shown in Figure 10, that is, a t distribution with location 0.350, scale 0.102, and degrees of freedom 3 is also the final result of the prior elicitation. Since the MATCH Uncertainty Elicitation Tool only allows fitting t distributions with three degrees of freedom, we also presented the expert with t distributions with different number of degrees of freedom using R (R Core Team, 2016). However, the expert decided that the t distribution with three degrees of freedom (i.e., $t(0.350, 0.102, 3)$) captured her prior knowledge best.

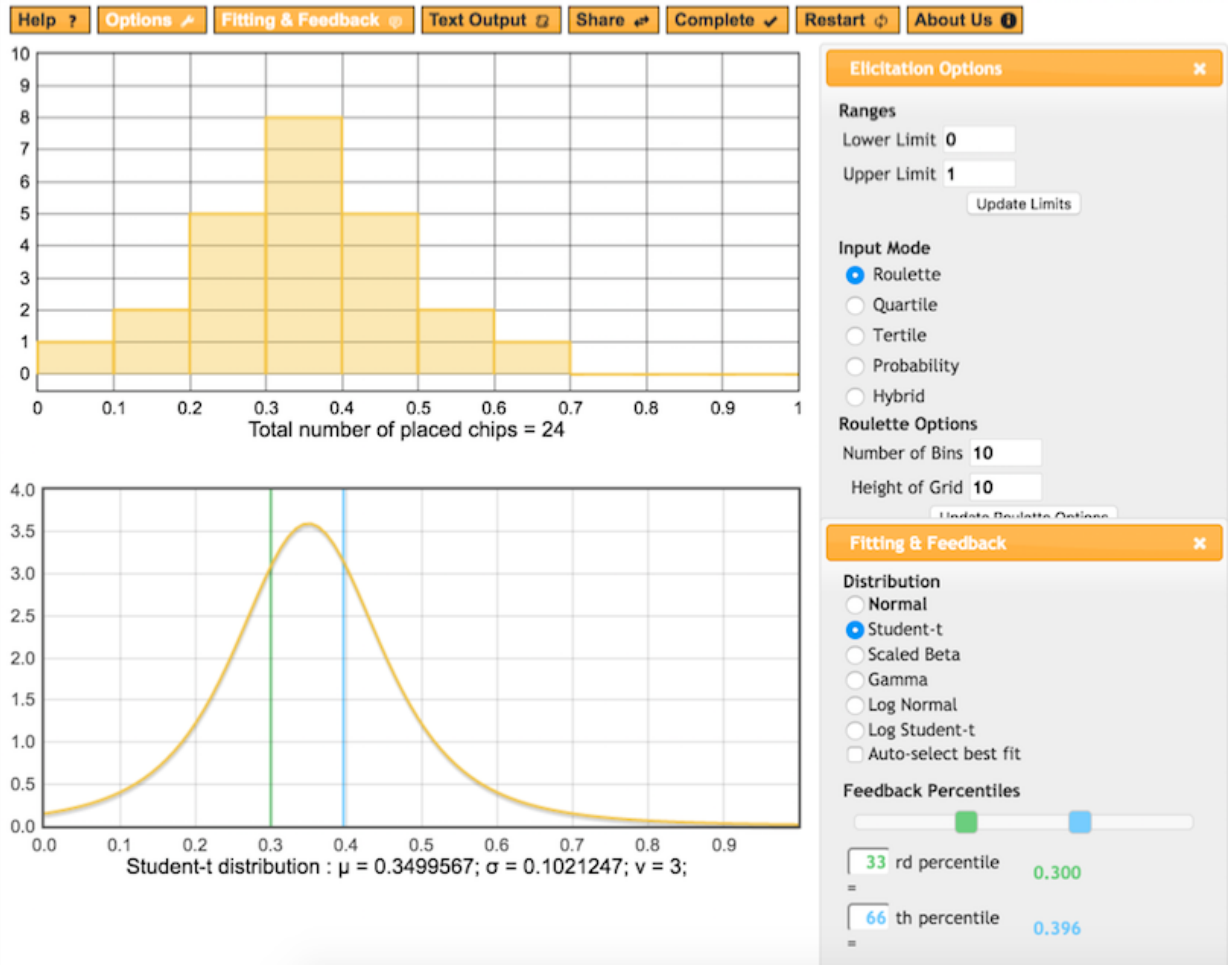

Figure 10: Depiction of the results of the prior elicitation using the MATCH Uncertainty Elicitation Tool (`http://optics.eee.nottingham.ac.uk/match/uncertainty.php`). The top part of the figure displays the results of the "Roulette" method where the expert has to allocate chips to bins in order to elicit an informed prior distribution, the bottom-part shows the fitted t distribution. The bottom-right corner displays the feedback that the expert received about the implied 33%-tile and 66%-tile values. Figure available at `https://tinyurl.com/jwqzt7c` under CC license `https://creativecommons.org/licenses/by/2.0/`.